

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide (translate or convert) relational query multidimensional

Searching within The ACM Digital Library for: (translate or convert) relational query multidimensi Found 28 of 256.817

REFINE YOUR SEARCH ▼ Refine by Keywords

Discovered Terms

▼ Refine by People Names

Institutions

▼ Refine by

Publications

Publication Year

Publication Names

ACM Publications

All Publications

Publishers

Content Formats

Authors Reviewers Search Results

Related Journals Related Magazines Related SI

Results 1 - 20 of 28

Sort by relevance

(translate or convert)

Save results to a Binder

Communications of the ACM: Volume 51 Issue 12

December 2008 Communications of the ACM

Publisher: ACM

Full text available: Digital Edition , Pdf (6.91 MB) Additional Information: full citation

Bibliometrics: Downloads (6 Weeks): 457. Downloads (12 Months): 4241. Cita

2 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ia Lefohn

August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes Publisher: ACM Propulation Request Permissions

Full text available: Pdf (63.03 MB) Additional Information: full citation, abstract, a

Bibliometrics: Downloads (6 Weeks): 137, Downloads (12 Months): 1493, Cita The graphics processor (GPU) on today's commodity video cards has ev-

▼ Refine by Conferences Sponsors Events

Proceeding Series

and flexible processor. The latest graphics architectures provide tremen computational horsepower, with fully programmable vertex ...

ADVANCED SEARCH Advanced Search

FEEDBACK

Please provide us with feedback

Found 28 of 256.817

Optimizing object queries using an effective calculus Leonidas Fegaras, David Maier

December 2000 Transactions on Database Systems (TODS), Volume 25 Is Publisher: ACM Prequest Permissions

Additional Information: full citation, abstract, r Full text available: Pdf (641.65 KB)

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 103, Citation

Object-oriented databases (OODBs) provide powerful data abstractions generally lack a suitable framework for query processing and optimization effective query optimizer is one of the key factors ...

Keywords: nested relations, object-oriented databases, query decorrel

4 Query processing techniques for arrays

Arunprasad P. Marathe, Kenneth Salem August 2002 The VLDB Journal - The International Journal on Very L Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (322.53 KB) Additional Information: full citation, abstract, r

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 58, Citation

Arrays are a common and important class of data. At present, database adequate array support: arrays can neither be easily defined nor conver array manipulations are not optimized. This paper describes ...

Keywords: Array manipulation language, Array query optimization, De Memory-usage optimization, Pipelined evaluation, User-defined function

 Modeling and querying multidimensional data sources in Siebel Anal system

Kazi A. Zaman, Donovan A. Schneider

June 2005 SIGMOD '05: Proceedings of the 2005 ACM SIGMOD internation of data

Publisher: ACM Pequest Permissions

Full text available: Pdf (401.92 KB) Additional Information: full citation, abstract.

Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 59, Citation

Large organizations have a multitude of data sources across the enterpr value from all of them. While the majority of these data sources may be data warehouse, many business units have their ...

⁶ ArchIS: an XML-based approach to transaction-time temporal database Fusheng Wang, Carlo Zaniolo, Xin Zhou

November 2008 The VLDB Journal — The International Journal on Ver

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (1.50 MB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 12, Downloads (12 Months): 54, Citation

Effective support for temporal applications by database systems represe objective that is difficult to achieve since it requires an integrated soluti including (1) expressive temporal representations ...

Keywords: Temporal database, Temporal grouping, Temporal guery, X

7 Query by Excel

Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Aman Naimat, L Allison Waingold

August 2005 VLDB '05: Proceedings of the 31st international conference or Publisher: VLDB Endowment

Full text available: Pdf (243.43 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 9, Downloads (12 Months): 72, Citation

Spreadsheets, and MS Excel in particular, are established analysis tools interface, provide an easy to use computational model, and offer substanalysis. However, as opposed to RDBMS, spreadsheets ...

8

The state of the art in distributed query processing
Donald Kossmann
December 2000 Computing Surveys (CSUR), Volume 32 Issue 4
Publisher: ACM Pagaest Permissions
Full text available: Pdf (455.39 KB)
Additional Information: full citation, abstract, i

Distributed data processing is becoming a reality. Businesses want to do often must do it in order to stay competitive. While much of the infrastr processing is already there (e.g., modern network ...

Bibliometrics: Downloads (6 Weeks): 115. Downloads (12 Months): 1009. Citi

Keywords: caching, dient-server databases, database application systems information systems, economic models for query processing, middlewar execution, query optimization, replication, wrappers

9 Fast detection of communication patterns in distributed executions. Thomas Kunz, Michiel F. H. Seuren

November 1997 CASCON '97: Proceedings of the 1997 conference of the C

Publisher: IBM Press

Full text available: Pdf (4.21 MB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 38, Downloads (12 Months): 387, Citatic

Understanding distributed applications is a tedious and difficult task. Vistime diagrams are often used to obtain a better understanding of the exvisualization tool we use is Poet, an event ...

- 10 Shape-based retrieval and analysis of 3D models
- Thomas Funkhouser, Michael Kazhdan
 August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Prequest Permissions

Full text available: Pdf (12.56 MB) Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 54, Downloads (12 Months): 644, Citatic

Large repositories of 3D data are rapidly becoming available in several I molecular blology, and computer graphics. As the number of 3D models need for computer algorithms to help people find ...

- 11 External memory algorithms and data structures: dealing with massi

 Deffrey Scott Vitter
- June 2001 Computing Surveys (CSUR), Volume 33 Issue 2

Publisher: ACM Prequest Permissions

Full text available: Pdf (828.46 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 121, Downloads (12 Months): 976, Citat

Data sets in large applications are often too massive to fit completely in memory. The resulting input/output communication (or I/O) between fa external memory (such as disks) can be a major performance ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hash hierarchical memory, multidimensional access methods, multilevel memsecondary storage, sorting

12 Seeing, hearing, and touching: putting it all together

Brian Fisher, Sidney Fels, Karon MacLean, Tamara Munzner, Ronald Rensin August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Paguest Permissions
Full text available: Pdf (20.64 MB)

Additional Information: full citation, cited by

Bibliometrics: Downloads (6 Weeks): 142, Downloads (12 Months): 1797, Cita

- 13 Towards integrated and efficient scientific sensor data processing: a Ji Wu, Yongluan Zhou, Karl Aberer, Kian-Lee Tan
- March 2009 EDBT '09: Proceedings of the 12th International Conference of Technology: Advances in Database Technology

Publisher: ACM Request Permissions

Full text available: Pdf (484.05 KB) Additional Information: full citation, abstract, I

Bibliometrics: Downloads (6 Weeks): 21, Downloads (12 Months): 75, Citation

In this work, we focus on managing scientific environmental data, which collected from wireless sensors. In environmental science applications, I validated, interpolated, aligned and aggregated ...

- 14 Estimating query result sizes for proxy caching in scientific database
 Tanu Malik, Randal Burns, Nitesh V, Chawla, Alex Szalay
- November 2006 SC '06: Proceedings of the 2006 ACM/IEEE conference on \$

 Publisher: ACM

 Full text available:

 Html (2.23 KB),

 Pdf (225.84 KB) Additional information: full citative conference on the conferenc

terms

Bibliometrics: Downloads (6 Weeks): 13, Downloads (12 Months): 42, Citation

In a proxy cache for federations of scientific databases it is important to before making a caching decision. With accurate estimates, near-optima obtained. On the other extreme, inaccurate estimates ...

Keywords: data mining, proxy caching, scientific federations

- 15 Visualizing geospatial data
- Theresa Marie Rhyne, Alan MacEachren, Theresa-Marie Rhyne August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Prequest Permissions

Full text available: Pdf (14.01 MB)

Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 107, Downloads (12 Months): 984, Citat

This course reviews concepts and highlights new directions in GeoVisual integrating geospatial data and geographic information systems (GIS) w visualization (VIS) methods. These include: • ...

16

Secondary bitmap indexes with vertical and horizontal partitioning Guadalupe Canahuate, Tan Apaydin, Ahmet Sacan, Hakan Ferhatosmanogl March 2009 EDBT '09: Proceedings of the 12th International Conference or Technology: Advances in Database Technology

Publisher: ACM Request Permissions
Full text available: Pdf (2.78 MB)

Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 9, Downloads (12 Months): 54, Citation

Traditional bitmap indexes are utilized as a special type of primary or cl queries are answered by performing fast logical operations supported by mapped to the physical data by using the row id of each ...

17 Concise descriptions of subsets of structured sets

Alberto O. Mendelzon, Ken Q. Pu June 2003 PODS '03: Proceedings of the twenty-second ACM SIGMOD-SI Principles of database systems

Publisher: ACM Request Permissions

Full text available: Pdf (334.12 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 2, Downloads (12 Months): 19, Citation

We study the problem of economical representation of subsets of struct with a set cover. Given a structured set U, and a language L whose exproblem of Minimum Description ...

18 Bridging the gap between OLAP and SQL

Jens-Peter Dittrich, Donald Kossmann, Alexander Kreutz

August 2005 VLDB '05: Proceedings of the 31st International conference or Publisher: VLDB Endowment

Full text available: Pdf (409.18 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 89, Citation

In the last ten years, database vendors have invested heavily in order t features for decision support. Examples of functionality that has been a 7], spreadsheet computations [19], grouping ...

19 Using AutoMed metadata in data warehousing environments

Hao Fan, Alexandra Poulovassilis

November 2003 DOLAP '03: Proceedings of the 6th ACM international work

Publisher: ACM Prequest Permissions

Full text available: Pdf (271.41 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 10, Downloads (12 Months): 72, Citation

What kind of metadata can be used for expressing the multiplicity of da transformation and integration processes in data warehousing environm further used for supporting other data warehouse activities? ...

Keywords: data integration, data warehouse, metadata

Alloy: a lightweight object modelling notation Daniel Jackson

April 2002 Transactions on Software Engineering and Methodology Publisher: ACM • Request Permissions

Full text available: Pdf (346.87 KB) Additional Information: full citation, abstract, i

Bibliometrics: Downloads (6 Weeks): 30, Downloads (12 Months): 220, Citatic

Alloy is a little language for describing structural properties. It offers a c with graphical object models, and a set-based formula syntax powerful constraints and yet amenable to a fully automatic ...

Keywords: Object models, Z specification language, first-order logic

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2009 A

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide (translate or convert) relational query multidimensional

37.000

Searching within The ACM Digital Library for: (translate or convert) relational guery multidimensi search)

Found 15 of 256,817

Refinements (remove all) click each refinement below to remove

Publication Year: 1992 ... 2004

Search Results Related Journals Related SIGs Related Confere Results 1 - 15 of 15 Sort by relevance in.

REFINE YOUR SEARCH ▼ Refine by Keywords

(translate or convert)

Discovered Terms

▼ Refine by People Names

Institutions Authors Reviewers

▼ Refine by Publications

Publication Year **Publication Names ACM Publications** All Publications Publishers

▼ Refine by Conferences Sponsors

Proceeding Series ADVANCED SEARCH

Advanced Search

FEEDBACK

Events

Please provide us with feedback

Found 15 of 256.817

Save results to a Binder

GPGPU; general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ia Woolley, Aaron Lefohn

August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Prequest Permissions

Full text available: Pdf (63.03 MB) Additional Information: full citation, abstract, cited to Bibliometrics: Downloads (6 Weeks); 137, Downloads (12 Months); 1493, Cita

The graphics processor (GPU) on today's commodity video cards has evextremely powerful and flexible processor. The latest graphics architecti tremendous memory bandwidth and computational horsepower, with fu vertex ...

2 Optimizing object queries using an effective calculus

Leonidas Fegaras, David Maier

December 2000 Transactions on Database Systems (TODS), Volume 25 I Publisher: ACM Propest Permissions

Additional Information; full citation, abstract, refere Full text available: Pdf (641.65 KB)

Bibliometrics: Downloads (6 Weeks): 6. Downloads (12 Months): 103. Citation

Object-oriented databases (OODBs) provide powerful data abstractions facilities, but they generally lack a suitable framework for query process optimization. The development of an effective query optimizer is one of

Keywords: nested relations, object-oriented databases, query decorrel optimization

3 Query processing techniques for arrays

Arunprasad P. Marathe, Kenneth Salem

August 2002 The VLDB Journal - The International Journal on Very L Volume 11 Issue 1

Publisher: Springer-Verlag New York, Inc.

Additional Information: full citation, abstract, refere Full text available: Pdf (322.53 KB)

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 58, Citation

Arrays are a common and important class of data. At present, database provide adequate array support: arrays can neither be easily defined no manipulated. Further, array manipulations are not optimized. This pape

Keywords: Array manipulation language, Array query optimization, De language, Memory-usage optimization, Pipelined evaluation, User-define

4 The state of the art in distributed query processing.

Donald Kossmann
December 2000 Computing Surveys (CSUR), Volume 32 Issue 4

Publisher: ACM Pequest Permissions

Full text available: Pdf (455.39 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 115, Downloads (12 Months): 1009, Cita

Distributed data processing is becoming a reality. Businesses want to do reasons, and they often must do it in order to stay competitive. While n infrastructure for distributed data processing is already there (e.g., moc

Keywords: caching, client-server databases, database application systematical systems. based information systems, economic models for query processing, mid architectures, query execution, query optimization, replication, wrapper

5 Fast detection of communication patterns in distributed executions Thomas Kunz, Michiel F. H. Seuren

November 1997 CASCON '97: Proceedings of the 1997 conference of the Co Studies on Collaborative research

Publisher: IBM Press

Full text available: Pdf (4.21 MB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 38. Downloads (12 Months): 387. Citatic

Understanding distributed applications is a tedious and difficult task. Vis on process-time diagrams are often used to obtain a better understandi of the application. The visualization tool we use is Poet, an event ...

- 6 Shape-based retrieval and analysis of 3D models
- Thomas Funkhouser, Michael Kazhdan

August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Preguest Permissions

Full text available: Pdf (12.56 MB) Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 54. Downloads (12 Months): 644. Citatic

Large repositories of 3D data are rapidly becoming available in several f mechanical CAD, molecular biology, and computer graphics. As the nurr grows, there is an increasing need for computer algorithms to help peor

- 7 External memory algorithms and data structures: dealing with massi. Jeffrey Scott Vitter
 - June 2001 Computing Surveys (CSUR), Volume 33 Issue 2

Publisher: ACM Propuest Permissions

Full text available: Pdf (828.46 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 121, Downloads (12 Months): 976, Cital

Data sets in large applications are often too massive to fit completely in internal memory. The resulting input/output communication (or I/O) be memory and slower external memory (such as disks) can be a major pe

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hash memory, hierarchical memory, multidimensional access methods, multil online, out-of-core, secondary storage, sorting

- 8 Seeing, hearing, and touching: putting it all together
- Brian Fisher, Sidney Fels, Karon MacLean, Tamara Munzner, Ronald Rensin August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Request Permissions

Full text available: Pdf (20.64 MB) Additional Information: full citation, cited by

Bibliometrics: Downloads (6 Weeks): 142, Downloads (12 Months): 1797, Cita

- 9 Visualizing geospatial data
- Theresa Marie Rhyne, Alan MacEachren, Theresa-Marie Rhyne August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Request Permissions

Full text available: Pdf (14.01 MB) Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 107. Downloads (12 Months): 984. Citat

This course reviews concepts and highlights new directions in GeoVisual four levels of integrating geospatial data and geographic information sy: scientific and information visualization (VIS) methods. These include: . .

- 10 Concise descriptions of subsets of structured sets
- Alberto O. Mendelzon, Ken Q. Pu

June 2003 PODS '03: Proceedings of the twenty-second ACM SIGMOD-SIG symposium on Principles of database systems

Publisher: ACM Request Permissions

Full text available: Rdf (334.12 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 2, Downloads (12 Months): 19. Citation

We study the problem of economical representation of subsets of struct sets equipped with a set cover. Given a structured set U, and a language expressions define subsets of U, the problem of Minimum Description ...

- 11 Using AutoMed metadata in data warehousing environments
- Hao Fan, Alexandra Poulovassilis November 2003 DOLAP '03: Proceedings of the 6th ACM international work warehousing and OLAP

Publisher: ACM Request Permissions

Full text available: Pdf (271.41 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 10. Downloads (12 Months): 72. Citation

What kind of metadata can be used for expressing the multiplicity of da data transformation and integration processes in data warehousing envi this metadata be further used for supporting other data warehouse activ

Keywords: data integration, data warehouse, metadata

12 Alloy: a lightweight object modelling notation

Daniel Jackson

April 2002 Transactions on Software Engineering and Methodology (Issue 2

Publisher: ACM Request Permissions

Full text available: Pdf (346.87 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 30. Downloads (12 Months): 220. Citatic

Alloy is a little language for describing structural properties. It offers a c compatible with graphical object models, and a set-based formula synta to express complex constraints and yet amenable to a fully automatic ...

Keywords: Object models, Z specification language, first-order logic

13 PIROL: a case study for multidimensional separation of concerns in:

engineering environments Stephan Herrmann, Mira Mezini

October 2000 OOPSLA '00: Proceedings of the 15th ACM SIGPLAN conferen oriented programming, systems, languages, and applications

Publisher: ACM ARequest Permissions

Full text available: Pdf (441.79 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 7, Downloads (12 Months): 41, Citation

In this paper, we present our experience with applying multidimensiona concerns to a software engineering environment. By comparing two diff system, we show the importance of separating integration issues from t implementation ...

Keywords: component integration, domain-specific language, separat software engineering environment

Also published in:

October 2000 SIGPLAN Notices Volume 35 Issue 10

14 A visual interface technique for exploring OLAP data with coordinate

hierarchies

Mark Sifer

November 2003 CIKM '03: Proceedings of the twelfth international conferer

and knowledge management

Publisher: ACM Request Permissions

Full text available: Pdf (272.82 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 11, Downloads (12 Months): 48, Citation

Multi-dimensional data occurs in many domains while a wide variety of I visual interfaces for querying such data exists. But many of these interf applicable to OLAP, as they do not support use of dimension hierarchies

Keywords: OLAP, data exploration, hierarchies, interface

15 Rule-based optimization and query processing in an extensible geon system

Ludger Becker, Ralf Hartmut Güting

June 1992 Transactions on Database Systems (TODS), Volume 17 Issue 2
Publisher: ACM Preguest Permissions

Full text available: Pdf (3.35 MB)

Additional Information: full citation, abstract, refere terms, review

Bibliometrics: Downloads (6 Weeks): 9, Downloads (12 Months): 48, Citation

Gral is an extensible database system, based on the formal concept of a relational algebra. Many-sorted algebra is used to define any application its query execution language, and its optimiztion rules. In this paper we

Keywords: extensibility, geometric query processing, many-sorted algorielational algebra, rule-based optimization

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2009 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Playe



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

sal query translation mulitidimensional query

Searching within The Guide for: sql query translation mulitidimensional (start a new search) Found 140 of 1.340.015

▼ Refine by Keywords

sql query translation n

Discovered Terms

▼ Refine by People Names Institutions Authors Reviewers

▼ Refine by Publications Publication Year Publication Names ACM Publications All Publications Content Formats Publishers

▼ Refine by Conferences Sponsors Events Proceeding Series

ADVANCED SEARCH Advanced Search

FEEDBACK

Please provide us with feedback

Found 140 of 1,340,015

Search Results REFINE YOUR SEARCH Related Conferences

Related Journals Related Magazines

Sort by relevance

Related SI in in

Results 1 - 20 of 140 Save results to a Binder

Result page: 1 2 3 4 5 6

1 Modeling and guerying multidimensional data sources in Siebel Anal relational system

Kazi A. Zaman, Donovan A. Schneider

June 2005 SIGMOD '05: Proceedings of the 2005 ACM SIGMOD Internation Management of data

Publisher: ACM Pequest Permissions

Full text available: Pdf (401.92 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 59, Citation

Large organizations have a multitude of data sources across the enterpr obtain business value from all of them. While the majority of these data consolidated in an enterprise data warehouse, many business units have

2 Relevance measures for XML information retrieval Olli Luoma

June 2007 International Journal of Web and Grid Services . Volume 3 Iss Publisher: Inderscience Publishers

Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks); n/a. Downloads (12 Months); n/a. Citatic

In recent years, a lot of work has been carried out to develop efficient n and querying XML data. Most of the proposals have approached the subdatabase point of view, i.e., they have primarily aimed at providing exa

Keywords: XML data, XPath query language, information retrieval, rele semistructured documents

3 Proceedings of the 7th International Workshop on the Web and Data

with ACM SIGMOD/PODS 2004; colocated with ACM SIGMOD/POE Luis Gravano, Sihem Amer-Yahia

June 2004 WebDB '04: Proceedings of the 7th International Workshop on Databases: colocated with ACM SIGMOD/PODS 2004

Publisher: ACM

Additional Information: full citation

Bibliometrics: Downloads (6 Weeks): n/a. Downloads (12 Months): n/a. Citatic

Answering queries using views: A survey

Alon Y. Halevy

December 2001 The VLDB Journal — The International Journal on Ver Bases , Volume 10 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (308.74 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 23, Downloads (12 Months): 230, Citatic

The problem of answering queries using views is to find efficient method guery using a set of previously defined materialized views over the data accessing the database relations. The problem has recently received sig

Keywords: Data integration, Date warehousing, Materialized views, Ou Survey, Web-site management

- 5 The state of the art in distributed query processing

Donald Kossmann
December 2000 Computing Surveys (CSUR), Volume 32 Issue 4

Publisher: ACM Preguest Permissions

Full text available: Pdf (455.39 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 115, Downloads (12 Months): 1009, Cita

Distributed data processing is becoming a reality. Businesses want to do reasons, and they often must do it in order to stay competitive. While n infrastructure for distributed data processing is already there (e.g., moc

Keywords: caching, client-server databases, database application systematical systems. based information systems, economic models for query processing, mid architectures, query execution, query optimization, replication, wrapper

Form-based proxy caching for database-backed web sites: keywords Oiong Lup, Jeffrey F. Naughton, Wenwei Xue

May 2008 The VLDB Journal - The International Journal on Very Lai Volume 17 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (483.35 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 9. Downloads (12 Months): 104. Citation

Web caching proxy servers are essential for improving web performance and recent research has focused on making proxy caching work for data sites. In this paper, we explore a new proxy caching framework that exi

Keywords: Database-backed web sites, Web proxy caching

7 Semantic caching of Web queries Boris Chidlovskii, Uwe M. Borghoff March 2000 The VLDB Journal — The International Journal on Very La Volume 9 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (235.09 KB)

Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks); 4. Downloads (12 Months); 56. Citation

In meta-searchers accessing distributed Web-based information repositing a major issue. Efficient query processing requires an appropriate cach unfortunately, standard page-based as well as tuple-based caching medical page-based page-

Keywords: Experiments, Query algorithms, Region containment, Sema Signature files

- 8 Temporal statement modifiers
- Michael H., Böhlen, Christian S., Jensen, Richard Thomas Snodgrass
 December 2000 Transactions on Database Systems (TODS), volume 25 I:
 Publisher: ACM & Request Permissions

Full text available: Pdf (317.23 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 13, Downloads (12 Months): 96, Citation

A wide range of database applications manage time-varying data. Many languages have been proposed, each one the result of many carefully minteracting design decisions. In this article we advocate a different appr

Keywords: ATSQL, statement modifiers, temporal databases

- 9 SchemaSQL: An extension to SQL for multidatabase interoperability Laks V. S. Lakshmanan, Fereidoon Sadri, Subbu N. Subramanian
 - December 2001 Transactions on Database Systems (TODS), Volume 26 I:
 Publisher: ACM Request Permissions

Full text available: Pdf (435.89 KB) Additional Information: full citation, abstract, refere terms, review

Bibliometrics: Downloads (6 Weeks): 18, Downloads (12 Months): 146, Citatic We provide a principled extension of SQL, called SchemaSQL, that offer

uniform manipulation of data and schema in relational multidatabase sy a precise syntax and semantics of *SchemaSQL* in a manner that ...

Keywords: Information integration, SchemaSQL, multidatabase system views, schematic heterogeneity

- 10 Rewriting queries with arbitrary aggregation functions using views Sara Cohen, Werner Nutt, Yehoshua Sagiv
- June 2006 Transactions on Database Systems (TODS), volume 31 Issue 2
 Publisher: ACM ♦ Request Permissions

Full text available: Pdf (294.01 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks); 6. Downloads (12 Months); 83. Citation

The problem of rewriting aggregate queries using views is studied for compart aggregation functions and built-in predicates. Two types coviews are introduced for rewriting aggregate queries: pure candidates.

Keywords: View usability, guery equivalence, guery rewriting

11 Optimizing object queries using an effective calculus

Leonidas Fegaras, David Maier

December 2000 Transactions on Database Systems (TODS), Volume 25 I-Publisher: ACM Request Permissions

Full text available: Pdf (641.65 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 103, Citation

Object-oriented databases (OODBs) provide powerful data abstractions facilities, but they generally lack a suitable framework for query process optimization. The development of an effective query optimizer is one of

Keywords: nested relations, object-oriented databases, query decorrel optimization

12 Query by Excel

Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Aman Naimat, L Subramanian, Allison Waingold

August 2005 VLDB '05: Proceedings of the 31st international conference or bases

Publisher: VLDB Endowment

Full text available: Pdf (243.43 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 9. Downloads (12 Months): 72. Citation

Spreadsheets, and MS Excel in particular, are established analysis tools attractive user interface, provide an easy to use computational model, a interactivity for what-if analysis. However, as opposed to RDBMS, sprea

13 A constraint-based querying system for exploratory pattern discovery Francesco Bonchi, Fosca Giannotti, Claudio Lucchese, Salvatore Orlando, R Roberto Trasarti

March 2009 Information Systems , Volume 34 Issue 1

Publisher: Elsevier Science Ltd.

Additional Information: full citation, abstract, references, index terms

Bibliometrics: Downloads (6 Weeks): n/a. Downloads (12 Months): n/a. Citatic

In this article we present ConQueSt, a constraint-based querying syster the intrinsically exploratory (i.e., human-guided, interactive and iterativ discovery. Following the inductive database vision, our framework ...

Keywords: Constrained pattern mining, Data mining query languages, systems, Inductive databases, Interactive data mining

14 On querying geospatial and georeferenced metadata resources in G Zehua Liu, Ee-Peng Lim, Wee-Keong Ng, Dion H. Goh

May 2003 JCDL '03: Proceedings of the 3rd ACM/IEEE-CS joint conference Publisher: IEEE Computer Society

Full text available: Pdf (92.05 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 20, Citation

G-Portal is a web portal system providing a range of digital library servigeospatial and georeferenced resources on the Web. Among them are ti query subsystems that provide a central repository of metadata resourc

15 ArchIS: an XML-based approach to transaction-time temporal databa Fusheng Wang, Carlo Zaniolo, Xin Zhou

November 2008 The VLDB Journal — The International Journal on Ver Bases, Volume 17 Issue 6
Publisher: Springer-Verlaa New York, Inc.

Full text available: 120 MB)

Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 12, Downloads (12 Months): 54, Citation

Effective support for temporal applications by database systems repress technical objective that is difficult to achieve since it requires an integra several problems, including (i) expressive temporal representations ...

terms

Keywords: Temporal database, Temporal grouping, Temporal query, X XOuery

16 CubiST#: Evaluating Ad-Hoc CUBE Queries Using Statistics Trees

November 2003 Distributed and Parallel Databases, Volume 14 Issue 3
Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): n/a. Downloads (12 Months): n/a. Citatic

We report on a new, efficient encoding for the data cube, which results up of OLAP queries that aggregate along any combination of dimensions and categorical attributes. We are focusing on a class of queries called ϵ

Keywords: data cube, data warehouse, multi-dimensional OLAP, query statistics tree

17 Evaluating XML-extended OLAP queries based on a physical algebration Xuepeng Ym, Torben Bach Pedersen

November 2004 DOLAP '04: Proceedings of the 7th ACM international work warehousing and OLAP

Publisher: ACM Request Permissions

Full text available: Pdf (206.65 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 50, Citation

In today's OLAP systems, integrating fast changing data, e.g., stock qur a cube is complex and time consuming. The widespread use of XML mal that this data is available in XML format on the WWW. Thus, making XM

Keywords: OLAP, XML, data integration, physical algebra, guery semail

18 Flexible and efficient IR using array databases

Roberto Cornacchia, Sándor Héman, Marcin Zukowski, Arjen P. Vries, Petei January 2008 The VLDB Journal — The International Journal on Very Bases, Volume 17 Issus 1

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (364.32 KB) Additional Information: full citation, abstract, index

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 86, Citation

The Matrix Framework is a recent proposal by Information Retrieval (IR flexibly represent information retrieval models and concepts in a single array framework. We provide computational support for exactly this frai

Keywords: Array databases, Database compression, Information retrie optimization

19 Accelerating XPath evaluation in any RDBMS Torsten Grust, Maurice Van Keulen, Jens Teubner

March 2004 Transactions on Database Systems (TODS), Volume 29 Issue Publisher: ACM Request Permissions

Full text available: Pdf (781.01 KB) Additional Information: full citation, appendices and abstract, references, cited t

Bibliometrics: Downloads (6 Weeks): 18, Downloads (12 Months): 173, Citatic

This article is a proposal for a database index structure, the *XPath accel* been specifically designed to support the evaluation of XPath path expre the index is capable to support *all* XPath axes (including ...

Keywords: Main-memory databases, XML, XML indexing, XPath

20 Buffering databse operations for enhanced instruction cache perform ingren Zhou, Kenneth A. Ross

June 2004 SIGMOD '04: Proceedings of the 2004 ACM SIGMOD internation Management of data

Publisher: ACM Proguest Permissions

Full text available: Pdf (188.52 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 11. Downloads (12 Months): 86. Citation

As more and more query processing work can be done in main memory a significant cost component of database operations. Recent database r that most of the memory stalls are due to second-level cache data miss

Result page: 1 2 3 4 5 6

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2009 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Playe

Modeling and Querying Multidimensional Data Sources in Siebel Analytics: A Federated Relational System

Kazi A. Zaman Siebel Systems 2207 Bridgepointe Pkwy San Mateo CA 94404 1-650-477-1553

Kazi.Zaman@siebel.com

Donovan A. Schneider Siebel Systems 2207 Bridgepointe Pkwy San Mateo CA 94404 1-650-477-1517

Donovan Schneider@siebel.com

ABSTRACT

Large organizations have a multitude of data sources across the enterprise and want to obtain business value from all of them. While the majority of these data sources may be consolidated in an enterprise data warehouse, many busness units have their own data marts where analysis is carried out against data stored in multidimensional data structures. It is often critical to pose queries which span both these sources. This is a challenge since queries which span both these sources. This is a challenge since which span both these sources the differing models and query languages (SQL vs MDX). The Siebel Analytics Server enables this requirement to be fulfilled. In this paper, we describe how the multidimensional metadata is modeled relationally within Siebel Analytics, efficient SQL to MDX translation algorithms and the conversion protocols required to convert a multidimensional result into a relational rower.

1. INTRODUCTION

Most companies have made significant investments in data warehousing technologies and business intelligence tools to maximize the value of their enterprise wide information. While the majority of the data may be loaded and transformed into a central data warehouse, there still remain other data sources in the organization which are extremely valuable but are not in the warehouse.

The ability to query multiple federated data sources is key for answering critical business queries. For example, finance departments typically use multidimensional databases for budgeting that allow them to carry out sophisticated calculations that cannot be performed in a relational database. A common query compares budgets versus actuals where the actuals are stored in a data warehouse. Another example is real time reporting where a single report may need up to the minute headcount data from the HR system along with the latest information on the sales pipeline from the front office systems.

Permission to make digital or hard copies of all or part of this work for personal or classroom use in granted without fee provided that copies to not made or distributed for profit or commercial advantage and that copies to the commercial advantage and that copies commercial advantage and that copies or republish, to post on servers or to redistribute to lists, requires prior sweet for remission and/or a fee.

SIGMOD 2005, June 14-16, 2005, Baltimore, Maryland, USA Copyright 2005 ACM 1-59593-060-4/05/06 \$5.00 The problem of federation has been addressed extensively by the database community both in general and with regard to support for specific data sources [1, 5, 6]. In this paper, we address the problem of integrating multidimensional data sources in a federated system. Multidimensional data sources typically have custom access methods with their own proprietary query languages and API's. Most of the vendors in this segment now support the XMLA standard [4], an open industry-standard web services interface designed for online analytical processing (OLAP). The existence of a common standard supported by a large number of vendors has made it viable to invest resources in devising algorithms and techniques for federating this class of data sources.

Federating multidimensional systems is challenging. These systems have a rich set of metadata that included dimensions with multiple hierarchies and measures with aggregation rules. This rich metadata does not have any direct analog in the relational world but still has to be used to pose meaningful queries against these data sources. There exists a powerful query language MDX which has a number of specialized operators over this metadata and returns results in the form of a multidimensional dataset rather than a rowset. In this pager, we describe how these issues are addressed in the context of the Siebel Analytics Server, a federated relational system.

The structure of the paper is as described below. In Section 2, we discuss the architecture of the Analytics Server with an emphasis on query processing and data modeling. We define the scope of the multidimensional data source federation problem. In Section 3, we provide an overview of multidimensional data sources, the MDX query language and the XMLA protocol. In Section 4, we describe how multidimensional metadata is modeled within Siebel Analytics. In Section 5, we address how to generate MDX queries from combination of an internal query plan representation and multidimensional data source specific metadata. We describe how the multidimensional result set is converted back into rowsets. Section 6 summarizes the contributions of this paper and discusses areas of future work.

2. ARCHITECTURE OF THE SIEBEL ANALYTICS SERVER

The Siebel Analytics Server is an ANSI SQL compliant query server at the core of Siebel's Analytics/Business Intelligence Suite. The Siebel Analytics Suite has a web based front end (Siebel Analytics Web) which issues SQL queries to the Analytics Server via the DDBC protocol. These queries are posed against logical tables, a concept we describe below. The logical queries are translated into the appropriate queries against physical back end databases. The results from the backend databases are combined and further manipulated before being returned to the client. In summary, the Analytics Server is a Federated relational system with the capability of query execution but without a storage subsystem.

The SQL queries sent to the Analytics Server are posed against a logical business model. A business model presents business information in a manner that parallels business analysts' understanding of the business structure rather than the physical structure in which the data may be stored. The business model contains dimensions, hierarchies, levels, measures and other data warehousing concepts. Note that unlike multidimensional data sources, the internals are purely relational; these objects serve as an aid to modeline.

A key concept required to support business models is the logical trable. A logical table consisting of one or more logical columns is an abstraction above a physical table. It can be a subset of a physical table (a subset of columns or a subset of rows); it can combine the contents of two or more physical tables or it can be derived from other logical tables. Typically, logical tables reduce complexity in the information model because a single logical table can map to multiple physical tables. Similarly, a logical column is can abstraction above a physical column. It can be mapped to one or more physical columns, to a scalar expression involving physical columns to other logical columns.

Logical tables and columns can be mapped to multiple sources. These sources can originate from multiple databases of the same or different types. This federation capability supports horizontal and vertical fragmentation across data sources. The ability to specify alternative data sources allows users to model replicated data sources as well as support aggregation navigation by specifying alternative sources at different levels of granularity. The relationships between logical tables are specified in terms of joins. These joins may be foreign key joins based on logical keys specified for each logical table, or outer joins. The repository has metadata pertaining to the physical data sources being mapped. This includes connectivity in formation, information about the type of the data source and a features table describing the capabilities of the data source

During query processing the Analytics Server takes a SQL query as input, parses it and converts it into internal data structures. Based on the metadata information, it deduces which physical tables correspond to the logical tables referred to in the query and produces an initial query plan in terms of these tables. All aggregate navigation is carried out at this point. This initial query plan is then optimized by a rule based query compiler. The query compiler carries out optimizations like pushing aggregations below joins as well as transformations required to convert the plan to a state where we can generate SOL to ship to one or more backend databases. This includes determining which operations should be executed remotely and which operations should be executed within the Analytics Server. The goal is to execute as much functionality as possible at the back end database. This has the obvious advantage of transporting as small a set of data from the back end database(s) to the Analytics Server which is typically a dominant cost in query processing. Post processing steps are carried out on the results of the query and the results are then shipped back to the client.

We now examine the areas impacted by the requirement to support multidimensional data sources. Multidimensional data sources are modeled in the physical layer of the metadata repository as a new database type and are mapped to the business model layer in an identical fastion to relational databases. Since users construct queries against the business model, the physical data sources are completely transparent to them. The joins between multidimensional data sources and relational data sources are specified in the business model layer.

Supporting multidimensional data sources requires specialized metadata for hierarchies and levels to enable us to generate MDX rather than SQL. Since the structure of a query plan is relational, in essence our multidimensional code generation module is a SQL to MDX translator. Our design goal is to ship as broad a class of queries as possible — the ability to push down GROUP BY's being especially important. This satisfies the goal of this project which is the ability to efficiently combine data from relational and multidimensional data sources.

3. MULTIDIMENSIONAL DATA SOURCES, MDX AND THE XMLA PROTOCOL

Multidimensional data sources typically present the user with a dimensional view of data. The data is typically organized around two key concepts: dimensions and messures. Each dimension can have one or more hierarchies. For example the Time dimension may have a Year-> Quarter-> Month->Day hierarchy and/or a Year-> Weah interarchy. Measures are metrics that are of business interest like profit or sales. Multidimensional data sources are organized as cuebse which consist of a collection of dimensions of dimensions. Any useful business questions are answered by obtaining the value of one or more measures for selected combinations of dimensional values. This form of data modeling is typical in data warehousing secantics [2].

Multidimensional data providers often have their own storage subsystems which are optimized for efficient access of dimensional queries. These storage subsystems use extensive pre-calculation of aggregates at various levels to provide efficient performance. Many vendor products use the term OLAP (Online Anabytical Processing) to describe this space. The terms MOLAP and ROLAP are used to distinguish between the storage mechanism used: ROLAP products use relational storage while MOLAP products use multidimensional data structures. HOLAP systems combine aspects of both

One barrier to wider adoption of multidimensional data providers has been the absence of a standard interface for access. The has been the absence of a standard interface for access. The DISCOVER and EXECUTE. The DISCOVER and EXECUTE The DISCOVER and EXECUTE The DISCOVER and EXECUTE of the DISCOVER and exact provided and a standard or a standard and a standard and

The MDX language is specialized for querying multidimensional data. MDX has rich functionality for manipulating members and sets of members of a dimension by means of functions like Ancestors/g and Descendants/g. Specialized constructs for Time Series calculations exist. Some business calculations require different aggregation rules to be carried on different dimensions; MDX has the concept of a solve order for dealing with dimensional formula precedence. Advanced MDX concepts are described in detail in [3], we will focus on the subset of MDX that the Analytics Server generates.

A key difference between a MDX query and a SQL query is that a MDX query terms a multidimensional grid of cells as output to a query. The number of dimensions in this grid depends upon the number of dises specified in the query. Each axis of a query is composed of a set of tuples each of which can have one or more dimensions. A zero axis query will still return a single cell as output to the query. The basic structure of a MDX query is as shown below.

```
SELECT [<axis_specification>
     [, <axis_specification>...]]
FROM [<cube_specification>]
[WHERE [<sticer_specification>)]
```

In the template above, the cube specification is the name of the coube being queried. This is a single cube name and unlike SQL there is no syntax for specifying joins. The WHERE clause is optimal and is referred to as the slicer axis. The MDX WHERE clause is semantically distinct from its usage in SQL, it serves the clause is semantically distinct from its usage in SQL, it serves the SELECT clause to values specified in a tuple. We make this clearer by means of an example.

```
SELECT { [Measures].[Profit] } ON COLUMNS, { [Time].[2002], [Time].[2003] } ON ROWS FROM Sales WHERE ( [Store].[USA].[California] }
```

This query requests Profit for stores in California for the years 2002 and 2003. The output is a grid with two cells, one containing the profit for 2002 and the other for 2003. The sticer axis ensures that the profit is calculated only for California. We have not explicitly specified the aggregation rule, this is defined in the cube metadata. Note that there is syntax for explicitly referring to dimensional members. While in this query there is a two dimensional grid as output, a greater number of axes can be specified.

4. MODELING MULTIDIMENSIONAL DATA SOURCES IN SIEBEL ANALYTICS METADATA

We now describe how to model multidimensional data sources in Siebel Analytics. We map a multidimensional data source to a cubetable in the physical level of the metadata repository. This can be carried out automatically by using the import functionality of the metadata administration tool which obtains the relevant metadata from the snecified XMLA provider. The import functionalty uses the DISCOVER call of the XMLA protocol to map the metadata obtained into the Sichel Analytics primitives described below. A cubetable is a special version of a relational table augmented with special metadata that allows us to generate MDX queries against the datasource. In all other respects it behaves identically to a relational table. Cubetables are mapped to the logical layer of the repository in the same manner as relational tables. During query processing, MDX will be generated against cubetables rather than SQL.

Each cubetable is based on a single multidimensional cube. A cubetable consists of cube columns. A cube column behaves in the same function as a regular column but is annotated with special metadata (hierarchy and level information) used only for MDX code generation purposes. A cube column can either be a measure or a dimensional column. Cube Columns which are measures are amontated with an aggregation rule (For example, sales may be associated with aggregation rule SUM. This information is normally available for non-derived aggregation rules via the XMLA protocol). Dimensional columns are either level keys or properties. A level key uniquely denotes a level of a hierarchy while a property is functionally dependent upon a level key. We illustrate this by means of an example.

Consider a cube with two hierarchies: Time: Year-> Quarter -> Month and Geography; State -> City -> Store Name. Store Name has the property Store Manager. We have two measures Profit and Sales which both have aggregation rule SUM associated with them. We wish to map this cube into a cube table T.

The resultant cube table T has the following cube columns (Year, Quater, Month, State, City, Store Nanne, Manager, Profit, Sales). Profit and Sales are labeled as measures and have the aggregation rule SUM associated with them. The cube table T has two hierarchies Time and Geography associated with it. Time has the following levels: Year, Quater and Month. Each of these levels has a level key of the same name associated with it. The Geography heartery has 3 levels: State, City and Store Name. Each level has a level key of the same name associated with it. The Geography heartery has 3 levels: State, City and Store Name. Each level has sore hamper associated with it as a property. Cube Table T has a primary key which is the set of all cube columns which are also level keys, In this example this comaists of the set (Year, Quarter, Month, State, City, Store Name). Both measures are functionally dependent upon the key.

We have some restrictions on how cubes can be mapped to cuberbles. A single cuberble can contain only one hierarchy per dimension. If a user needs to map additional hierarchies, an additional cuberble will have to be created. Additionally, we assume that the hierarchy is a homogeneous, balanced hierarchy. This implies that every node of every level in the hierarchy is of the same type and every branch of the hierarchy has the same number of levels. This framework does not model parent child hierarchies which are self joining hierarchies where the number of levels is not explicitly specified up front. A typical example of levels is not explicitly specified up front. A typical example of these hierarchies is to model employee-manager relationships.

5. MDX CODE GENERATION AND OUERY PROCESSING

We now examine the steps taken to process a query. We then outline the issues specific to queries that reference a multidimensional data source and describe how they are handled.

- The user poses a SQL query against the logical model built in the Siebel Analytics metadata. This is converted into a logical query representation.
- The navigator transforms the logical query into a plan comprised solely of objects from the physical backend data sources.
- The query compiler/optimizer refines the query plan. At the completion of this phase the query plan is ready for execution. Note that this plan is purely relational in acture.
- 4. The execution plan may reference multiple physical tables. The plan is marked appropriately with indications of where each operation will be executed. Multidatabase joins and other post processing operations are marked for execution in the middle tier.
- 5. In a relational setting, the data corresponding to the portion of a query plan marked for remote execution can be obtained by converting the required portion of the plan into a SQL query. This conversion is referred to as code generation and results in the appropriate vendor dialect of SQL being produced.
- After execution of the SQL queries, data is fetched to the Analytics Server. Post processing is carried out and the resulting rowset is returned to the client

The overall flow of query processing requires modification when a backend data source is multidimensional. In sep 5 we produce a SQL query from a relational query plan fragment. For multidimensional data sources MDX rather than SQL is generated. This problem is referred to as the MDX Code Generation problem for the remainder of this paper. In sep 6, we refer to post processing carried out in the server on rowsets. However, the response to a MDX query is a multidimensional data set (in XML) and not a rowset. An unpivoting protocol is required which specifies how to convert from a multidimensional data set to a rowset. We first examine a simple unpivoting protocol in detail before moving to the broader problem of MDX Code Generation.

We describe a simple protocol used for Microsoft Analysis Server. We restrict the class of MDX queries to two dimensions. This implies that we will have a two dimensional grid of result in the class of a set of tuples on columns. This imples on columns are described by a formula specified in the MDX difference of the class a unique that the class of the

Example: We describe how an intermediate rowset is constructed for the MDX query below.

```
SELECT
[[Measures].[Sales]] on CDLUMNS,
[Crossjoin([[Year].Members],
[[Products].[Soda].Members]) on ROWS
FROM [Sales]
```

The XML output file will contain the member [Measures].[Sales] as the column tuple. Let the tuples returned on rows be (1998, Pepsi), (1998,Coke), (1997,Pepsi), (1997,Coke). We construct a rowset by appending the appropriate cell values to each row tuple.

The methodology described in the previous paragraph enables us to transform as XML output file into a rowset. However, the MDX query generated may not have an exact correspondence with the execution plan. This could be for a number of reasons: (1) All MDX queries return a measure even for dimensional only queries where the default measure is returned. This additional column in the obtained rowset may need to be pruned. (2) We may require a different ordering of columns in the project list. In the scheme specified earlier, all measures columns will be in columns following dimensional members of the row tunles.

Correct ordering of columns can be ensured by using a simple protocol consisting of a list of integers. The list (2, 1.4) implies that the mappings between the intermediate rowset and the final rowest are as follows. Intermediate Column(IC) 2 maps to Fical Column(FC) 1, IC 1 maps to FC2 and IC 4 maps to FC3. Note that not every column in the finermediate rowset has to map to a column in the final rowset. Every MDX query generated would have a corresponding protocol list generated. Different data sources may have different mapping protocols. The more complex this protocol, the greater is the overhead of middle tier processing.

5.1 MDX Code Generation

MDX code generation aims to generate a query based on the portion of the query execution plan marked for execution at the multidimensional database. The query plan is relational, in essence we are trying to translate 8QL to MDX. We translate a subset of SQL corresponding to the template below.

```
SELECT cl, c2... aggr(m1), aggr(m2)
FROM Table
WHERE <conditions>
GROUP BY cl, c2...
Conditions>
```

This template is typical of queries where users are interested in obtaining metrics at a specified level of aggregation. We support other query templates but for ease of exposition we restrict ourselves to this primary use case. We examine various code generation strategies using queries against the table in the example given below. Each queries paginst the table in the example given below. Each queries playing that the table in the adaption of the present the MDX code generation. We then present the MDX code generation algorithm in its entirety.

Consider a cube, Sales, with two hierarchies: Time (Levels: Year, Quarter) and Store (Levels: Store Country, Store State) and one measure: Unit Sales. In the corresponding repository we have a cubetable T(Store Country, Store State, Year, Owarter. Unit

Sales). Let the aggregation rule associated with Unit Sales be SUM.

Query 1 [Multiple dimensions plus measure with matching aggregate rule]

Select Store Country, Year, SUM(Unit Sales)

```
From T3

Group By Store Country, Year

Select
[not Sales]] on columns,
[nonemptyrrossjoin [Store Country].members,
[Year].members]] on rows
```

This query requires us to crossjoin two dimensions to obtain the tuples to be placed on rows. We can directly use the member [Unit Sales] because the aggregation rule in the query matches that of the measure.

Query 2 [Measure with non-matching aggregate rule]

```
Select Store Country, Year, AVG(Unit Sales)
From T
Group By Store Country, Year
with
 set [A] as '{[Store Country].members}'
 set [B] as '{[Year].members]
 set [C] as 'nonemptycrossjoin({[A]},{[B]})'
 member (measures).[MS1] as
'AVG (nonemptycrossjoin (Descendants (Store, currentme
mber, [Store State]),
Descendants (Time.currentmember, [Quarter]) ), [Unit
Sales1)'
select
    ([MS1]) on columns.
    ([C]) on rows
from [Sales]
```

For each combination of country and year we need to find the average sales. The measure needs to be explicitly computed because AVG differs from the aggregation rule associated with [Unit Sales] which is SUM. This requires us to model the measure as a calculated measure [MSI], where each value at the grain of Store state and Quarter is examined to compute the average sale. We are making use of the MDX WITH section to define named sets and calculated measures. Sets A and B are named sets which are used to store the dimensional members while set C consists of the tuples obtained by crossjoning A and to obtain all possible combinations of countries and years.

Query 3 [Single dimension with multiple levels plus measure]

```
Select Store Country, Store State, SUM(Unit Sales)
From T
Group By Store Country, Store State
```

```
with
member [measures].[CountryAnc] as
'ancestor(Store.Currentmember,[Store
Country)).name'
```

```
set [A] as 'Descendants[[Store],[Store State])'
select
[[Measures].[CountryAnc],[Unit Sales]] or
columns,
[ [A]) on rows
From [Sales]
```

In this query we require multiple levels of the same dimension. In the MDX language a tuple can only contain one member from a single dimension. We handle this by modeling Store Country, the the velocities of the root, as a calculated member using the Ancestor of function. If we changulated member using the Ancestor of SELECT list we would still generate the same MDX query. This is an example of a query where the post processing protocol and the processing protocol control of the processing protocol control of the processing protocol.

Query 4 [Multiple dimensions plus measure with matching aggregate rule, predicate refers to GROUP BY columns]

Where Store Country In ('USA', 'India') AND Year = '1997'

Select Store Country, Year, SUM(Unit Sales)

Group By Store Country, Year

From T

```
Alternative I
with

set (A) as '[Store Country].members'
set (B) as '[Store].members'
set (B) as '[Store].members'
set (C) as 'nonemptyprossjoin({[A]},{[B]})'
select
{[Unit Sales]} on columns,
$Pileter([C]),[Store_currentmember.name = "USA"
OR Store_currentmember.name = "India") AND
$From [Sales]

Alternative 2
```

```
Alemative 2

With
set [A] as 'ffilter([Store Country].members,
Store.currentmember.name = "USA" OR
Store.currentmember.name = "India"))'
set [B] as 'ffilter([Year].members,
time.currentmember.name = "1997"),
set [C] as 'nonemptycrossjoin([A]], [B])'
select
[C] as 'nonemptycrossjoin([A]], [B])'
select
[C] for rows
```

This query illustrates how the WHERE clause is processed. In the first MDX query in WHERE clause is exceuted after the crossjoin. In the second query we break the WHERE clause into constituents which are applied on each dimension before the crossjoin. The second query is more efficient since a crossjoin is a very expensive operation and its input sets are smaller in this case. Relational systems typically have optimizers that are responsible for optimizations like push down of predicates. Optimizers for suitidimensional data sources are less mature and hence it is critical to generate the second query. This optimization cannot be carried out for all classes of predicates. For example, if we substitute the AND in the query with an OR the disjunctive predicate with have to be evaluated after the crossjoin.

Query 5 [Multiple levels of multiple dimensions plus measure with matching aggregate rule, predicates refer to unprojected and projected levels]

```
Select Store Country, Store State, year, SUM(Unit Sales)
From T
Where Store Country = 'USA' AND Store State = 'CA' AND
quarter IN ('Q3', 'Q4')
Group By Store Country, Store State, Year
```

```
member [Measures].[CountryAnc] as
   'ancestor (Store . Currentmember, [Store
Country]).name'
set [A] as 'filter( {[Store Country
| members | Store currentmember name = "USA") '
set[B] as'Filter( Generate ({[A]}, Descendants
([Store].currentmember, [Store].[Store State])),
[Store].currentmember.name= "CA" )'
set [C] as '{[Time].[Year].members}'
        as 'nonemptycrossjoin({[B]],{[C]})'
member (measures), [MSI] as
'SUM(filter(Descendants(Time.currentmember,[Quarte
r]), Time.currentmember.name = "Q3" OR
Time.currentmember.name = "Q4"),[Unit Sales])'
 select
       ([Measures].[CountryAnc],
```

```
{[Measures].[CountryAnc],
    [Measures].[MS1]] on columns,
    [[D]) on rows
From [Sales]
```

In this query, we have predicates on multiple levels of a dimension. We apply the predicate to the named set containing members at the Country level. We use the generate function to obtain the descendants at the state level only of the members of the previously created named set. Notice that we explicitly compute the sum of sales at the quarter grain. This illustrates that even for matching aggregate rules, we cannot always use precomputed aggregates at the backend data source.

5.1.1 MDX Code Generation Algorithm for SELECT-FROM-WHERE-GROUP BY Queries The Input: SQL Query

Output: MDX query and corresponding Protocol List

- Examine SELECT list of query. If multiple levels from a dimension are present, create the appropriate calculated A prestor members.
- Convert WHERE clause to CNF. If a conjunct can be associated with a single level of a single dimension, add it to the list of conjuncts marked for early evaluation.
- For every dimension, create a named set for each projected level. If any filters marked for early evaluation are applicable, apply them in the definition of the named set. If named sets are created for multiple levels of a dimension, the latest named set is sourced from previously created named sets.
- Create a named set ([Q]) corresponding to the set obtained by crossjoining the lowest projected named set for each dimension. Applicable cross dimensional filters are applied here.

- For each aggregate in the project list which does not have a matching aggregate rule, create a new calculated member which aggregates over all required cells.
- Output all ancestors, measures and calculated members on columns, named set [O] on rows.
- If HAVING clause is present, apply predicates on [O]

Other query templates such as SELECT-FROM-WHERE queries will have a code generation algorithm which is similar in spirit.

6. CONCLUSIONS AND FUTURE WORK

In this paper we described how to model and query multidimensional data sources within Siebel Analyties. This is achieved by modeling multidimensional data sources as cubetables - relational tables augmented with hierarchy and level metadata in the physical layer. We showed examples of various SQL querier stanslated to MDX and highlighted the performance issues with different alternatives. We examined how various classes of WHERE Cause preclicates are supported in MDX. Given that the result of a MDX query is a multidimensional data set, we described how this output is mapped to a relational rowset. This ability to handle both relational and multidimensional data in a single integrated framework fulfils a critical business need.

We are pursuing several areas of future work. There is considerable variance in MDX support from different vendors. Restricted MDX support implies modifying our MDX code generation algorithms and post processing protocosls. Overall query response time is best when we minimize the amount of data feethed to the middle tier. We aim to support as wide a template of SQL queries as possible so that the resultant MDX results in a minimal set of data being fetched. Cubetables currently model only balanced hierarchies. Modeling parent child hierarchies needs new metadata and new MDX code generation algorithms.

7. REFERENCES

- V. Josifovski, P. Schwarz, L. Haas and E. Lin, Garlie: a new flavor of federated query processing for DB2. Proceedings of ACM SIGMOD 2002., 524-532.
- [2] R. Kimball. The Data Warehouse Toolkit: The Complete Guide to Dimensional Modeling. John Wiley & Sons, 2002.
- [3] G. Spofford. MDX Solutions. Wiley, 2001.
- [4] http://www.xmla.org
- [5] I. Manolescu, D. Florescu, D. Kossmann. Answering XML Queries on Heterogeneous Data Sources. *Proceedings of VLDB* 2001., 241-250
- [6] D. Chamberlin, J. Robie, D. Florescu. Quitt: An XML Query Language for Heterogeneous Data Sources. Proceedings of WebDB 2001., 1-25



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide sql guery translation mulitidimensional virtual table

114.4

Searching within **The Guide** for: sql query translation mulitidimensional virtual table (start a new s Found 20 of 1,340,015

Refinements (remove all) click each refinement below to remove

Publication Year: 1992 ... 2004

REFINE YOUR SEARCH Search Results

Refine by Keywords Results 1 - 20 of 20

sql query translation n

Discovered Terms

▼ Refine by People

Names Institutions Authors Reviewers

Refine by Publications Publication Year Publication Names ACM Publications All Publications Content Formats Publishers

▼ Refine by Conferences Sponsors Events Proceeding Series

ADVANCED SEARCH

Advanced Search

FEEDBACK

Please provide us with feedback

Found 20 of 1,340,015

Related Journals Related SIGs Related Confere Sort by relevance

Save results to a Binder

1 The state of the art in distributed query processing

Donald Kossmann

December 2000 Computing Surveys (CSUR), Volume 32 Issue 4

Publisher: ACM Pequest Permissions

Full text available: Pdf (455.39 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 115, Downloads (12 Months): 1009, Cita

Distributed data processing is becoming a reality. Businesses want to do reasons, and they often must do it in order to stay competitive. While n infrastructure for distributed data processing is already there (e.g., mod

Keywords: caching, client-server databases, database application systbased information systems, economic models for query processing, mid architectures, query execution, query optimization, replication, wrapper

2 Answering queries using views: A survey

Alon Y. Halevy

December 2001 The VLDB Journal — The International Journal on Ver

Bases , Volume 10 Issue 4
Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (308.74 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 23. Downloads (12 Months): 230. Citatic

The problem of answering queries using views is to find efficient methor query using a set of previously defined materialized views over the data accessing the database relations. The problem has recently received sig

Keywords: Data integration, Date warehousing, Materialized views, Qu Survey, Web-site management

3 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ia Woolley, Aaron Lefohn August 2004 **STGGRAPH '04:** SIGGRAPH 2004 Course Notes Publisher: ACM Pequest Permissions

Full text available: Pdf (63.03 MB) Additional Information: full citation, abstract, cited in

Bibliometrics: Downloads (6 Weeks): 137, Downloads (12 Months): 1493, Cite

The graphics processor (GPU) on today's commodity video cards has evextremely powerful and flexible processor. The latest graphics architecti tremendous memory bandwidth and computational horsepower, with fuvertex.

4 Evaluating XML-extended OLAP queries based on a physical algebr

Xuepeng Yin, Torben Bach Pedersen November 2004 DOLAP '04: Proceedings of the 7th ACM international work warehousing and OLAP

Publisher: ACM Prequest Permissions

Full text available: Pdf (206.65 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 50, Citation

In today's OLAP systems, integrating fast changing data, e.g., stock que a cube is complex and time consuming. The widespread use of XML mal that this data is available in XML format on the WWW. Thus, making XM

Keywords: OLAP, XML, data integration, physical algebra, guery semail

MIL primitives for querying a fragmented world Peter A. Boncz, Martin L. Kersten

October 1999 The VLDB Journal — The International Journal on Very Bases , Volume 8 Issue 2

Publisher: Springer-Verlag New York. Inc.

Full text available: Pdf (261.36 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 37, Citation

In query-intensive database application areas, like decision support and systems that use vertical fragmentation have a significant performance order to support relational or object oriented applications on too of such

Keywords: Database systems, Main-memory techniques, Query langua optimization, Vertical fragmentation

6 Accelerating XPath evaluation in any RDBMS

Torsten Grust, Maurice Van Keulen, Jens Teubner

March 2004 Transactions on Database Systems (TODS), Volume 29 Issue Publisher: ACM Request Permissions

Full text available: Pdf (781.01 KB) Additional Information: full citation, appendices and abstract, references, cited t

Bibliometrics: Downloads (6 Weeks): 18, Downloads (12 Months): 173, Citatic

This article is a proposal for a database index structure, the XPath accel been specifically designed to support the evaluation of XPath path expre

the index is capable to support all XPath axes (including ...

Keywords: Main-memory databases, XML, XML indexing, XPath

7 The Design and Implementation of a Corporate Householding Knowl to Improve Data Quality

Stuart Madnick, Richard Wang, Xiang Xian

December 2003 Journal of Management Information Systems, Volume 2 Publisher: M. E. Sharpe, Inc.

Additional Information: full citation, abstract, references, cited by

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

Advances in corporate householding are needed to address certain catequality problems caused by data misinterpretation. In this paper, we fire of these data quality problems and our more recent results from studyir

Keywords: Context Mediation, Corporate Household, Corporate Housel Quality, Database Interoperability, Enterprise Knowledge Management

8 GridDB: a data-centric overlay for scientific grids David T. Liu, Michael J. Franklin

August 2004 VLDB '04: Proceedings of the Thirtieth International con large data bases - Volume 30 , Volume 30

Publisher: VLDB Endowment

Full text available: Pdf (300.42 KB) Additional Information; full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 3. Downloads (12 Months): 13. Citation

We present GridDB, a data-centric overlay for scientific grid data analys currently deployed process-centric middleware, GridDB manages data e processes. GridDB provides a suite of services important to data analysi

- 9 Model-driven development of Web applications: the AutoWeb system Piero Fraternali, Paolo Paolini
 - October 2000 Transactions on Information Systems (TOIS), Volume 18 I Publisher: ACM Request Permissions

Additional Information: full citation, abstract, refere Full text available: Pdf (6.94 MB)

Bibliometrics: Downloads (6 Weeks): 49. Downloads (12 Months): 425. Citatic

This paper describes a methodology for the development of WWW applienvironment specifically tailored for the methodology. The methodology development environment are based upon models and techniques alreahypermedia, ...

Keywords: HTML, WWW, application, development, intranet, modeling

10 Specifying OLAP Cubes on XML Data Mikael R. Jensen, Thomas H. Møller, Torben Bach Pedersen December 2001 Journal of Intelligent Information Systems, Volume 17 In Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

On-Line Analytical Processing (OLAP) enables analysts to gain insight at fast and interactive access to a variety of possible views on information dimensional model. The demand for data integration is rapidly becoming

Keywords: OLAP, XML, data integration, data warehousing, multidimer

11 Knowledge Mining With VxInsight: Discovery ThroughInteraction George S. Davidson, Bruce Hendrickson, David K. Johnson, Charles E. Mey November 1998 Journal of Intelligent Information Systems, Volume 11 I Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

The explosive growth in the availability of information is overwhelming I information management systems. Although individual pieces of information find, the larger context in which they exist has become harder t

Keywords: browsing, graphical user interface, information retrieval, invisualization

12 Streams, structures, spaces, scenarios, societies (5s): A formal mod

Marcos André Gonçalves, Edward A. Fox, Layne T. Watson, Neill A. Kipp April 2004 Transactions on Information Systems (TOIS), volume 22 Issu-Publisher: ACM ❖ Reguest Permissions

Full text available: Pdf (316.85 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 44, Downloads (12 Months): 396, Citatic

Digital libraries (DLs) are complex information systems and therefore diffeoundations lest development efforts diverge and interoperability suffers propose the fundamental abstractions of Streams, Structures, Spaces, S

Keywords: applications., definitions, foundations, taxonomy

13 Fast detection of communication patterns in distributed executions. Thomas Kunz, Michiel F. H. Seuren

November 1997 CASCON '97: Proceedings of the 1997 conference of the Ci Studies on Collaborative research

Publisher: IBM Press

Full text available: Pdf (4.21 MB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 38. Downloads (12 Months): 387. Citation

Understanding distributed applications is a tedious and difficult task. Vis on process-time diagrams are often used to obtain a better understandi of the application. The visualization tool we use is Poet, an event ...

14 Effect of node size on the performance of cache-conscious B*-trees

Richard A. Hankins, Jignesh M. Patel

June 2003 SIGMETRICS '03: Proceedings of the 2003 ACM SIGMETRICS in conference on Measurement and modeling of computer systems

Publisher: ACM Request Permissions

Full text available: Pdf (271.16 KB) Additional Information: full citation, abstract, refere terms

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 74, Citation

In main-memory databases, the number of processor cache misses has the performance of the system. Cache-conscious indices are designed to performance by reducing the number of processor cache misses that an

Keywords: B*-tree, cache-conscious, index

Also published in:

June 2003 SIGMETRICS Performance Evaluation Review Volume 31 Issue 1

15 Survey of Spatio-Temporal Databases Tamas Abraham, John F. Roddick

March 1999 Geoinformatica, Volume 3 Issue 1

Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): n/a. Downloads (12 Months): n/a. Citatic

Spatio-temporal databases aim to support extensions to existing model: Information Systems (SIS) to include time in order to better describe or environment. Although interest into this area has increased in the past number ...

Keywords: spatio-temporal databases, survey

16 Modelling and Manipulating Multidimensional Data in Semistructured Raymond K. Wong, Franky Lam, M. A. Orgun

October 2001 World Wide Web , Volume 4 Issue 1-2

Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

Multidimensional information is pervasive in many computer application series, spatial information, data warehousing, and visual data. While ser or XML is becoming more and more popular for information integration.

Keywords: XML, multidimensional data, multidimensional logics, semis datahases

- 17 Using AutoMed metadata in data warehousing environments
- Hao Fan, Alexandra Poulovassilis

November 2003 **DOLAP '03:** Proceedings of the 6th ACM international work warehousing and OLAP

Publisher: ACM Paguest Permissions

Full text available: Pdf (271.41 KB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 10. Downloads (12 Months): 72. Citation

What kind of metadata can be used for expressing the multiplicity of da data transformation and integration processes in data warehousing envi this metadata be further used for supporting other data warehouse activities.

Keywords: data integration, data warehouse, metadata

18 Toward improved geographic information services within a digital goof the NSE digital government initiative geographic information syste Louis Hecht, Barbara Kucera

May 2000 dg.o '00: Proceedings of the 2000 annual national conference or government research

Publisher: Digital Government Research Center

Full text available: Pdf (531.35 KB) Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 9, Downloads (12 Months): 109, Citation

This material is based upon work supported in part by the National Scie under Grant No. EIA-9818131. Any opinions, findings, and conclusions are recommendations expressed in this material are those of the authors are necessarily reflect ...

- 19 Query evaluation techniques for large databases
- Goetz Graefe

June 1993 Computing Surveys (CSUR), Volume 25 Issue 2

Publisher: ACM Request Permissions

Full text available: Pdf (9.37 MB)

Additional Information: full citation, abstract, refere terms, review

Bibliometrics: Downloads (6 Weeks): 118, Downloads (12 Months): 995, Citat

Database management systems will continue to manage large data volualgorithms for accessing and manipulating large sets and sequences will provide acceptable performance. The advent of object-oriented and extr

Keywords: complex query evaluation plans, dynamic query evaluation database systems, iterators, object-oriented database systems, operato parallelization, parallel algorithms, relational database systems, set-mal sort-hash duality

- 20 A declarative approach to optimize bulk loading into databases
- Silhem Amer-Yahia, Sophie Cluet
 June 2004 Transactions on Database Systems (TODS), Volume 29 Issue 2
 Publisher: ACM PRoduest Permissions

Full text available: Pdf (1.00 MB) Additional Information: full citation, abstract, refere

Bibliometrics: Downloads (6 Weeks): 10. Downloads (12 Months): 123, Citatio

Applications, such as warehouse maintenance, need to load large data in The efficiency of loading depends on the resources that are available at the target systems. Our work aims to understand the performance crite

Keywords: Declarative bulk loading, algebra, recovery, side-effects

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2009 ACM, Inc Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Playe



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide sql query translation conversion mulitidimensional

000000

in in

Searching within The Guide for: sql query translation conversion mulitidimensional (start a new se Found 24 of 1,340,015

Refinements (remove all) click each refinement below to remove

Publication Year: 1992 ... 2004

REFINE YOUR SEARCH

Search Results

Related Journals Related SIGs Related Confere Sort by relevance

Results 1 - 20 of 24 Save results to a Binder

Result page:

sal query translation c

▼ Refine by Keywords

Discovered Terms

▼ Refine by People Names Institutions Authors Reviewers

▼ Refine by Publications Publication Year Publication Names ACM Publications All Publications Content Formats Publishers

▼ Refine by Conferences Sponsors Events Proceeding Series

ADVANCED SEARCH

Advanced Search

FEEDBACK

Please provide us with feedback

Found 24 of 1.340.015

Temporal statement modifiers

Michael H. Böhlen, Christian S. Jensen, Richard Thomas Snodgrass December 2000 Transactions on Database Systems (TODS), Volume 25 I

Publisher: ACM Prequest Permissions Full text available: Pdf (317.23 KB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 13, Downloads (12 Months): 96, Citation

A wide range of database applications manage time-varying data. Many languages have been proposed, each one the result of many carefully m interacting design decisions. In this article we advocate a different appr

Keywords: ATSOL, statement modifiers, temporal databases

SchemaSQL: An extension to SQL for multidatabase interoperability

Laks V. S. Lakshmanan, Fereidoon Sadri, Subbu N. Subramanian
December 2001 Transactions on Database Systems (TODS), Volume 26 1 Publisher: ACM Prequest Permissions

Additional Information: full citation, abstract, referei Full text available: Pdf (435.89 KB)

Bibliometrics: Downloads (6 Weeks): 18, Downloads (12 Months): 146, Citatic

We provide a principled extension of SOL, called SchemaSOL, that offer uniform manipulation of data and schema in relational multidatabase sy a precise syntax and semantics of SchemaSQL in a manner that ...

Keywords: Information integration, SchemaSQL, multidatabase system views, schematic heterogeneity

Optimizing object queries using an effective calculus.

Leonidas Fegaras, David Maier

December 2000 Transactions on Database Systems (TODS), Volume 25 In Publisher: ACM Request Permissions

Full text available: Ptt (641.65 KB) Additional Information: full citation, abstract, referen terms, review

Bibliometrics: Downloads (6 Weeks); 6. Downloads (12 Months); 103. Citation

Object-oriented databases (OODBs) provide powerful data abstractions facilities, but they generally lack a suitable framework for query process optimization. The development of an effective query optimizer is one of

Keywords: nested relations, object-oriented databases, query decorrel optimization

4 CubiST**: Evaluating Ad-Hoc CUBE Queries Using Statistics Trees Joachim Hammer, Lixin Fu

November 2003 Distributed and Parallel Databases , Volume 14 Issue 3 Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

We report on a new, efficient encoding for the data cube, which results up of OLAP queries that aggregate along any combination of dimensions and categorical attributes. We are focusing on a class of gueries called a

Keywords: data cube, data warehouse, multi-dimensional OLAP, query statistics tree

- 5 Evaluating XML-extended OLAP queries based on a physical algebra Xuepeng Yin, Torben Bach Pedersen
- November 2004 DOLAP '04: Proceedings of the 7th ACM international work warehousing and OLAP

Publisher: ACM ARequest Permissions

Full text available: Pdf (206.65 KB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 3. Downloads (12 Months): 50. Citation

In today's OLAP systems, integrating fast changing data, e.g., stock gue into a cube is complex and time consuming. The widespread use of XML possible that this data is available in XML format on the WWW. Thus, m

Keywords: OLAP, XML, data integration, physical algebra, guery semail

- 6 GPGPU: general purpose computation on graphics hardware
- David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ia Woolley, Aaron Lefohn

August 2004 SIGGRAPH '04: SIGGRAPH 2004 Course Notes

Publisher: ACM Request Permissions

Full text available: Pdf (63.03 MB) Additional Information; full citation, abstract, cited to

Bibliometrics: Downloads (6 Weeks): 137, Downloads (12 Months): 1493, Cita

The graphics processor (GPU) on today's commodity video cards has ev extremely powerful and flexible processor. The latest graphics architecti tremendous memory bandwidth and computational horsepower, with fuvertex ...

7 MIL primitives for querying a fragmented world

Peter A. Boncz, Martin L. Kersten

October 1999 The VLDB Journal — The International Journal on Very Bases . Volume 8 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (261.36 KB) Additional Information: full citation, abstract, reference terms

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 37, Citation

In query-intensive database application areas, like decision support and systems that use vertical fragmentation have a significant performance order to support relational or object oriented applications on top of such

Keywords: Database systems, Main-memory techniques, Query langua optimization, Vertical fragmentation

8 The Design and Implementation of a Corporate Householding Knowl to Improve Data Quality

Stuart Madnick, Richard Wang, Xiang Xian

December 2003 Journal of Management Information Systems, volume 2 Publisher: M. E. Sharpe, Inc.

Additional Information: full citation, abstract, references, cited by

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

Advances in corporate householding are needed to address certain cate quality problems caused by data misinterpretation. In this paper, we fir of these data quality problems and our more recent results from studyir

Keywords: Context Mediation, Corporate Household, Corporate Housel Quality, Database Interoperability, Enterprise Knowledge Management

Query-based data warehousing tool
 Rami Rifaleh, Nabila Aicha Benharkat

November 2002 **DOLAP '02:** Proceedings of the 5th ACM international work

Publisher: ACM Pequest Permissions

Full text available: Pdf (280.49 KB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 34, Downloads (12 Months): 313, Citatic

Data warehousing is an essential element of decision support. It aims at knowledge user to make better and faster daily business decisions. In decisional database, meta-data is needed to enable the communication

Keywords: data warehouse, mapping expression, meta-data, query ba

Model-driven development of Web applications: the AutoWeb system

Piero Fraternali, Paolo Paolini

October 2000 Transactions on Information Systems (TOIS), Volume 18 I

Publisher: ACM Pequest Permissions

Full text available: Pdf (6.94 MB)

Additional Information: full citation, abstract, referenterms

Bibliometrics: Downloads (6 Weeks): 49, Downloads (12 Months): 425, Citatic

This paper describes a methodology for the development of WWW applienvironment specifically tailored for the methodology. The methodology development environment are based upon models and techniques alrea hypermedia. ...

Keywords: HTML, WWW, application, development, intranet, modeling

11 Streams, structures, spaces, scenarios, societies (5s): A formal mod

libraries

Marcos André Gonçalves, Edward A. Fox, Layne T. Watson, Neili A. Klpp April 2004 Transactions on Information Systems (TOIS), volume 22 Issur Publisher: ACM ♦ Request Permissions

Full text available: Pdf (316.85 KB) Additional Information: full citation, abstract, referenterms, review

Bibliometrics: Downloads (6 Weeks): 44, Downloads (12 Months): 396, Citatic

Digital libraries (DLs) are complex information systems and therefore defoundations lest development efforts diverge and interoperability suffers we propose the fundamental abstractions of Streams, Structures, Space

Keywords: applications., definitions, foundations, taxonomy

12 A Distributed Geographic Information System on the Common Objec Architecture (CORBA)

Fangju Wang March 2000 **Geoinformatica** , Volume 4 Issue 1

Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a. Citatic

A distributed geographic information system (distributed GIS) has stren reliability, efficiency, resource sharing, and flexibility for incremental sy: However, developing a distributed GIS is a challenging task. The major

Keywords: distributed information system, geographic information sys processing, query optimization, the Common Object Request Broker Arc

13 Integrating symbolic images into a multimedia database system usin and abstraction approaches

Aya Soffer, Hanan Samet
December 1998 The VLDB Journal — The International Journal on Ver
Bases , Volume 7 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: Pdf (227.30 KB) Additional Information: full citation, abstract, references

Bibliometrics: Downloads (6 Weeks): 20, Downloads (12 Months): 94, Citation

Symbolic images are composed of a finite set of symbols that have a se Examples of symbolic images include maps (where the semantic meanir is given in the legend), engineering drawings, and floor plans. Two appr

Keywords: Image indexing, Multimedia databases, Query optimization content. Spatial databases. Symbolic-image databases

14 Fast detection of communication patterns in distributed executions Thomas Kunz, Michiel F. H. Seuren

November 1997 CASCON '97: Proceedings of the 1997 conference of the Conference of the Conference of Collaborative research

Publisher: IBM Press

Full text available: Pdf (4.21 MB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 38, Downloads (12 Months): 387, Citatic

Understanding distributed applications is a tedious and difficult task. Vis on process-time diagrams are often used to obtain a better understandi of the application. The visualization tool we use is Poet, an event ...

- 15 Complexity and expressive power of logic programming
- Evgeny Dantsin, Thomas Eiter, Georg Gottlob, Andrei Voronkov September 2001 Computing Surveys (CSUR), Volume 33 Issue 3

Publisher: ACM Pequest Permissions

Full text available: Pdf (552,99 KB) Additional Information: full citation, abstract, referenterms

Bibliometrics: Downloads (6 Weeks): 40, Downloads (12 Months): 371, Citatic

This article surveys various complexity and expressiveness results on di logic programming. The main focus is on decidable forms of logic progra particular, propositional logic programming and datalog, but we also me

Keywords: Complexity, datalog, expressive power, logic programming, logic, query languages

- 16 A foundation for representing and guerying moving objects
- Ralf Hartmut Güting, Michael H. Böhlen, Martin Erwig, Christian S. Jensen, Lorentzos, Markus Schneider, Michalis Vazirgiannis

March 2000 Transactions on Database Systems (TODS), Volume 25 Issue Publisher: ACM Request Permissions

Full text available: Pdf (268.05 KB) Additional Information: full citation, abstract, referentierms

Bibliometrics: Downloads (6 Weeks): 20, Downloads (12 Months): 161, Citatic

Spatio-temporal databases deal with geometries changing over time. The

is to provide a DBMS data model and query language capable of handlir dependent geometries, including those changing continuously that desc

Keywords: abstract data types, algebra, moving objects, moving point spatio-temporal data types, spatio-temporal databases

17 Survey of Spatio-Temporal Databases

Tamas Abraham, John F. Roddick

March 1999 Geoinformatica, Volume 3 Issue 1

Publisher: Kluwer Academic Publishers

Full text available: Publisher Site Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citatic

Spatio-temporal databases aim to support extensions to existing model: Information Systems (SIS) to include time in order to better describe or environment. Although interest into this area has increased in the past number ...

Keywords: spatio-temporal databases, survey

18 Toward improved geographic information services within a digital goof the NSF digital government initiative geographic information syste Louis Hecht, Barbara Kucera

May 2000 dg.o '00: Proceedings of the 2000 annual national conference or government research

Publisher: Digital Government Research Center

Full text available: Pdf (531.35 KB) Additional Information: full citation, abstract

Bibliometrics: Downloads (6 Weeks): 9, Downloads (12 Months): 109, Citation

This material is based upon work supported in part by the National Scie under Grant No. EIA-9818131. Any opinions, findings, and conclusions $\mathfrak c$ recommendations expressed in this material are those of the authors ar necessarily reflect ...

19 Analytical view of business data

Adam Yeh, Jonathan Tang, Youxuan Jin, Sam Skrivan

August 2004 KDD '04: Proceedings of the tenth ACM SIGKDD international Knowledge discovery and data mining

Publisher: ACM Pequest Permissions

Full text available: Pdf (337.00 KB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 21, Downloads (12 Months): 156, Citatic

This paper describes a logical extension to Microsoft Business Framewor Analytical New (AV). AV consists of three components: Model Service ft Business Intelligence Entity (BIE) for programming model, and IntellDri

Keywords: OLAP, OLTP, analytics, application framework, business interpersistence, information retrieval and navigation, object-oriented

- 20 Using AutoMed metadata in data warehousing environments
- Hao Fan, Alexandra Poulovassilis November 2003 DOLAP '03: Proceedings of the 6th ACM international work warehousing and OLAP

Publisher: ACM Pequest Permissions

Full text available: Pdf (271.41 KB) Additional Information: full citation, abstract, referen

Bibliometrics: Downloads (6 Weeks): 10, Downloads (12 Months): 72, Citation

What kind of metadata can be used for expressing the multiplicity of da data transformation and integration processes in data warehousing envi can this metadata be further used for supporting other data warehouse

Keywords: data integration, data warehouse, metadata

Result page:

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2009 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



USPTO September 03, 2009

mber 03, 20

My Account

Manage Account Purchased Disclosures

Publish ...

A New Disclosure

Search By ...

Full Text Concept Document ID

. ...

View ...

My Recent Activities/Events My Published Disclosures Most Recently Published

Navigate To ...

Main Page Support Logout No records matched your search.

Perhaps you should try a less restrictive query.

Search query: relational query (translation or conversion) multidimensional q
Published Before: 12-4-2004 (Original publication date)

New search | Modify this search



September 03, 2009

My Account

Manage Account Purchased Disclosures

Publish ...

A New Disclosure

Search By ...

Full Text Concept Document ID

View ...

My Recent Activities/Events My Published Disclosures Most Recently Published

Navigate To ...

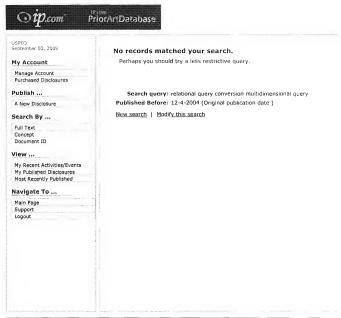
Main Page Support Logout No records matched your search.

Perhaps you should try a less restrictive query.

Search query: relational query translation multidimensional query

Published Before: 12-4-2004 (Original publication date)

New search | Modify this search







September 03, 2009

My Account

Manage Account Purchased Disclosures

Publish ...

A New Disclosure

Search By ...

Full Text Concept Document ID

View ...

My Recent Activities/Events My Published Disclosures Most Recently Published

Navigate To ...

Main Page Support Logout

No records matched your search.

Perhaps you should try a less restrictive query.

Search query: relational multidimensional query Published Before: 12-4-2004 (Original publication date)

New search | Modify this search



USPTO September 03, 2009

My Account

Manage Account Purchased Disclosures

Publish ...

A New Disclosure

Search By ...

Full Text Concept Document ID

View ...

My Recent Activities/Events My Published Disclosures Most Recently Published

Navigate To ...

Main Page Support Logout No records matched your search.

Perhaps you should try a less restrictive query.

Search guery: relational guery translation

Published Before: 12-4-2004 (Original publication date)

New search | Modify this search



USPTO September 03, 2009

My Account

Manage Account Purchased Disclosures

Publish ...

A New Disclosure

Search By ...

Full Text Concept Document ID

View ...

My Recent Activities/Events My Published Disclosures Most Recently Published

Navigate To ...

Main Page Support Logout Displaying records #1 through 10 out of 36

Result # 1 Relevance: OCOCO

Extensible Design for Generating Alternative Join Sequences In a R 1990-04-01 IPCOM000100452D

This article relates to a flexible method for developing an optimal plan to execute a renumerates different orders in which the tables in a query may be joined and elimin-provide optimal ...

Result # 2 Relevance: OCOCO

Method for Determining All Predicates And Expressions Implied by Ouery

1990-03-01 IPCOM000100293D

This invention determines all predicates and expressions that are implied by a set of For example, the predicates p1:X.a \approx Y.b and p2:Y.b = Z.c together imply a predica expression e(..., ...

Result # 3 Relevance: OOO

Relational Query Optimization Using a Process Network

1978-05-31 IPCOM000150541D

Leland L. Beck Department of Computer Science Southern Methodist University Dalla translation of queries from a variety of relational lan- guages into a fundamental set illustrated with ...

Result # 4 Relevance: OOO

Adaptive Access Plan for Select Queries With Multiple Predicates 1990-01-01 IPCOM000099329D

A technique is described whereby an adaptive access plan algorithm is implemented selection predicates. The plan uses indexes involved to measure the number of tuple access path is chosen to minimize ...

Result # 5 Relevance: OO

RULE-BASED QUERY OPTIMIZATION IN EXTENSIBLE DATABASE SYS

RULEBASED QUE OPTIMIZATION IN EXTENSIBLE DATABASE SYSTEMS GOETZ GRAE of the requirements for the degree of Doctor of Philosophy (Computer Sciences) at t 1987 vi Table of ...

Result # 6 Relevance: 🔾

QUERY OPTIMIZATION FOR RELATIONAL DATABASE SYSTEM 1980-10-31 IPCOM000151741D

Report No. UIUCDCS-R-80-1034 QUEXY OPTIMIZATION FOR RELATIONAL DATABASI Department of Computer Science Cniversity of Illinois at Urbana-Champaign Urbana in part by the National Science ...

Result # 7 Relevance: OC

A DESIGN MEHTODOLOGY FOR A UNIVERSAL RELATION SCHEME IN 1982-12-31 IPCOM000151479D

A DESIGN METHODOLOGY FCIW A UNIVERSAL RELATION SCHEME IMPLEMENTATION S., Bowling Green State University, 1974 M. S., Bowling Green State University, 197

the requirements for the degree ...

Result # 8 Relevance: OO

Access Path Selection in Relational Database Systems

988-02-01 IPCOM000056992D

A technique is described whereby a query optimizer selects the least expensive acce in database management systems. The concept considers the distribution of attribut queries, It ...

Result # 9 Relevance: OO

Method for Compact Storage and Retrieval of Relational Data Bases Length Encoding

Length Encod

01 IPCOM000043063D

A completely inverted file capable of general relational access and storage of genera achieved with greatly reduced storage requirements by utilizing a Sparse Matrix Ass. Data Structuring and ...

Result # 10 Relevance: 🛇

PATTERN-BASED AND KNOWLEDGE-DIRECTED QUERY COMPILATION FOR RECURSI's submitted in partial: fulfillment of the requirements for the degree of Doctor of Philo UNIVERSITY OF WISCONSIN MADISON ...

Displaying page 1 of 4 << FIRST | < BACK | NEXT > | LAST >>

Search query: relational query

Published Before: 12-4-2004 (Original publication date)

New search | Modify this search | Search within current results

Someh



Home I Login I Logout I Access Information I Alerts I Purchase History I

Welcome United States Patent and Trademark Office

BROWSE

SEARCH

IEEE XPLORE GUIDE

Search Results

Results for "((relational query translation)<in>metadata)"

Your search matched 2 of 2392400 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.



» Search Options

View Session History

New Search

» Key

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

AIP JNL AIP Journal

AVS JNL AVS Journal

IEEE CNF IEEE Conference

Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard

Indicates open access

content

Modify Search

standards.

((relational query translation)<in>metadata)

Check to search only within this results set

IEEE/IET/AIP/AVS Books Educational Courses A
IEEE/IET journals, transactions, letters, magazines, conference proceedings, AIP/

view selected items | Select All Deselect All

1. A Generalized Implementation Method for Relational |
Sublanguages

Beck, L.L.:

Software Engineering, IEEE Transactions on

Volume SE-6, Issue 2, March 1980 Page(s):152 - 162

AbstractPlus | Full Text: PDF(2560 KB) | IEEE JNL Rights and Permissions

CONTRACTOR AND ADDRESS OF THE ADDRES

2. A theory of translation from relational queries to hiera

Weiyi Meng; Yu, C.; Won Kim; Knowledge and Data Engineering, IEEE Transactions on Volume 7, Issue 2, April 1995 Page(s):228 - 245 Digital Object Identifier 10.1109/69,382294

AbstractPlus | Full Text: PDF(1680 KB) IEEE JNL

Rights and Permissions

Help Contact Us

Copyright 20





Home | Login | Logout | Access Information | Alerts | Purchase History |

Welcome United States Patent and Trademark Office

13 Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "((relational query translation)<in>metadata) and multidimensional"

Your search matched 0 of 2392400 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.



» Search Options

View Session History

New Search

» Key

IEEE Journal or IEEE JNL Magazine

IET JNL IET Journal or Magazine

AIP JNL AIP Journal

AVS JNL AVS Journal

IEEE Conference

IEEE CNF Proceeding

IET CNF IET Conference

Proceeding

IEEE STD IEEE Standard

Indicates open access content

Modify Search

((relational query translation)<in>metadata) and multidimensional

Bearch.

Check to search only within this results set

Display Format: @ Citation C Citation & Abstract

IEEE/IET/AIP/AVS

Books

Educational Courses

IEEE/IET journals, transactions, letters, magazines, conference proceedings, AIP/

view selected items Select All Deselect All

No results were found.

standards.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan

Indeped by M inspec Help Contact Us

Copyright 20

Seegh.

IEEE XPLORE GUIDE



Home | Login | Logout | Access Information | Alerts | Purchase History |

SEARCH

Welcome United States Patent and Trademark Office

C Search Results BROWSE

Results for "((relational query translation multidimensional)<in>metadata)"

Your search matched 0 of 2392400 documents. A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.



» Search Options

View Session History

New Search

» Kev

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

AIP JNL AIP Journal

AVS JNL AVS Journal

IEEE CNF IEEE Conference

Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard content

Indicates open access

Modify Search

((relational query translation multidimensional)<in>metadata) Check to search only within this results set

Display Format: @ Citation C Citation & Abstract

IEEE/IET/AIP/AVS **Educational Courses**

Books IEEE/IET journals, transactions, letters, magazines, conference proceedings, AIP/ standards.

view selected items Select All Deselect All

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistan



Help Contact Us

Copyright 20

Saurch



Home | Login | Logout | Access Information | Alerts | Purchase History |

Welcome United States Patent and Trademark Office

BROWSE

SEARCH

IEEE XPLORE GUIDE

Educational Courses

Search Results

Results for "((relational query)<in>metadata) and multidimensional query"

Your search matched 1 of 2392400 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Refevance in Descending order.



» Search Options

View Session History

New Search

» Key

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

AIP JNL AIP Journal

AVS JNL AVS Journal

IEEE CNF IEEE Conference

Proceeding

IFT Conference

IET CNE Proceeding

IEEE STD IEEE Standard

Indicates open access

content

Modify Search

((relational query)<in>metadata) and multidimensional query

Check to search only within this results set

Display Format: @ Citation C Citation & Abstract

IEFE/IET/AIP/AVS Books

IEEE/IET journals, transactions, letters, magazines, conference proceedings, AIP/

standards.

view selected items Select All Deselect All

1. Supporting Multi-attribute Queries in Peer-to-Peer Da Systems

Min Yu: Zhanhuai Li: Longbo Zhang:

Parallel and Distributed Computing, Applications and Technologies, 2007. PC International Conference on

3-6 Dec. 2007 Page(s):515 - 522

Digital Object Identifier 10.1109/PDCAT.2007.66

AbstractPlus | Full Text: PDF(334 KB) IEEE CNF

Rights and Permissions

Help Contact Us

හ Copyright 20



Web Images Videos Maps News Shopping Gmail more -

Search settings | Sign in



query translation

Search Advanced Search

Web Show options...

Results 1

Query Translation

Boolean Query Translation is an ongoing project in the Stanford Digital Library Project. It deals with the problem of translating Boolean queries into. infolab.stanford.edu/~kevin/queryTranslation.html - Cached - Similar

PDFI Semantics Preserving SPARQL-to-SQL Query Translation for Optional ... File Format: PDF/Adobe Acrobat - View as HTML in this paper, we presented the SPARQL-to-SQL query translation under the ... to-SQL query translation algorithm, SPARQLtoSQL, for SPARQL queries that ...

www.cs.wayne.edu/~artem/main/research/TR-DB-052006-CLJF.pdf - Similar by A Chebolko - 2006 - Cited by 19 - Related articles

Query expansion and query translation as logical inference

A number of studies have examined the problems of query expansion in monolingual information Retrieval (IR), and query translation for crosslanguage IR. ... portal.acm.org/citation.cfm?id=767402 - Similar

by JY Nie - 2003 - Cited by 12 - Related articles - All 4 versions

CiteSeerX — A Query Translation Scheme for Rapid Implementation of ... We focus on the query translation component of the toolkit, called the converter. The converter takes as input a Query Description and Translation Language ... citeseer.ist.psu.edu/papakonstantinou95query.html - Cached - Similar by Y Papakonstantinou - 1995 - Cited by 231 - Related articles - All 12 aversions

[PDF] Query Translation by Text Categorization

File Format: PDF/Adobe Acrobat - View

Cuery Translation by Text Categorization. Patrick Ruch. SIM, University Hospital of Geneva. 24 Micheil du Crest. 1201 Geneva, Switzerland ... www.aclweb.org/anthology/C/O4/C04-1099.pdf - Similar

by P Ruch - Cited by 13 - Related articles - All 14 versions

[PDF] Efficient XML-to-SQL Query Translation: Where to Add the Intelligence? File Format: PDF/Adobe Acrobat - View as HTML.

or one could use a more intelligent query translation al- gorithm and attempt to generate ... this information during the runtime query translation. ...

www.cis.upenn.edu/~susan/cis650/AddIntelligence.pdf - Similar

by R Krishnamurthy - 2004 - Cited by 39 - Related articles - All 10 versions

MrCoM: A Cost Model for Range Query Translation in Deep Web Data ...
Using the MrCoM, the query translation strategy for new global range queries can be inferred. We also propose a Pre-processing-based Stepwise Algorithm ...
www2.computer.org/portal/web/csdl/doi/10.../SKG.2008.69 - Cached - Similar

[PDF] WikiTranslate: Query Translation for Cross-lingual Information ... File Format: PDF/Adobe Acrobat - View

Main approaches to query translation are dictionary based translation, the use of parallel ... possibilities of Wikipedla for query translation in CLIR. ... www.def-campaign.org/2008/...notes/nguyen-paperCLEF2008.pdf - Similar by D Nguyen - Otted by 1 - Related articles [PDF] From Query Translation to Cross-Language Query Expansion with ...
File Format: PDF/Adobe Acrobat - View as HTML

Dictionary-based approaches to query translation have been ... of query translation that combines other types of term relation to ...

research.microsoft.com/en.../cikm2007.mcmodel.camera_ready.pdf - Similar by G Cao - Related articles - All 5 versions

Approximate Query Translation (Extended Version) - Stanford ... Dec 27, 2008 ... Our results show that for query translation we need to handle interdependencies among both query conjuncts as well as disjuncts ...

ilpubs.stanford.edu:8090/385/ - Cached - Similar by C Chang - 1999 - Cited by 3 - Related articles - All 8 versions

1 2 3 4 5 6 7 8 9 10 Next

query translation

Sea

Search within results - Language Tools - Search Help - Dissatisfied? Help us im

Google Home - Advertising Programs - Business Solutions - Priva

Web Images Videos Maps News Shopping Gmail more -

Search settings I Sign in



query translation multidimensional

Search Advanced Search

Web Show options ...

Results 1 - 10 of about 12'

Scholarly articles for query translation multidimensional

Phrasal translation and query expansion techniques for ... - Ballesteros - Cited by 237 A query language for multidimensional arrays. Design. - Libkin - Cited by 91 nD-SQL: A Multi-dimensional Language for ... - Gingras - Cited by 44

Mapping and query translation between xml, objects, and relations
A request is received to perform a query of a multi-dimensional database, ... Receive info
on patent apps like Mapping and query translation between xml, ...
www.freshpatents.com/-d/20090611ptan20090150367.php - Cached - Similar

multidimensional query definition of multidimensional query in the ...
Encyclopedia article about multidimensional query. Information about multidimensional query in ... Multidimensional Systems · Multidimensional translation ...
encyclopedia2 thefreedictionary.com/multidimensional-query - Cached - Similar

Query costing in a multidimensional database - US Patent 7392242 ...
US Patent 7392242 - Query costing in a multidimensional database Query formulation, input preparation, or translation707/5Query augmenting and refining ...
www.patentstorm.us/patents/7392242/claims.html - Similar

Report then query capability for a multidimensional database model ...

US Patent 7222130 - Report then query capability for a multidimensional database model ... Query formulation, input preparation, or translation707/101, ...

www.patentstorm.us/patents/7222130/claims.html - Similar

Show more results from www.patentstorm.us

resi Query Translation Method for Cross Language Information Retrieval File Format. Adobe PostScript - View as HTML 3 introduces our basic idea for the query translation. across languages. ... we use a multidimensional vector space, called word ...

www-csli.stanford.edu/semlab/infomap/Papers/mtsummit-CLIR.ps - Similar by H Masuichi - Cited by 4 - Related articles - All 4 versions

Information Sciences: A grain preservation translation algorithm...
However, to the best of our knowledge, no algorithm has been proposed that can systematically translate an entire ER diagram into a multidimensional model...
linkinghub elsever.com/retrieve/pii/S0020025507001673 - Similar
by TC Chen - 2007 - Related articles

рог Multi-dimensional Annotation and Alignment in an English-German ... File Format. PDF/Adobe Acrobat - <u>View</u>

Since the multi-dimensional annotation and alignment is realised in XML, the queries are posed using XQuery, 5. This query language is ... www.aclweb.org/anthology/W/W06/W06-2705.pdf - Similar by S Hansen-Schirra - Cited by 14 - Related articles - All 10 versions

Extending Query Translation to Cross-Language Query Expansion with ... Dictionary-based approaches to query translation have been The optimization problem can be cast as the multi-dimensional ... portal.acm.org/ft_gateway.cfm?id=1321491&type=pdf - Similar by G Cao - 2007 - Related articles - All 5 versions

[PPT] YAM2 A Multidimensional Conceptual Model
File Format: Microsoft Powerpoint - View as HTML

Multidimensional Algebra. November 7th, 2003. Alberto Abelló. 13. Basic query. Translation to SQL. November 7th, 2003. Alberto Abelló ...

www.cis.drexel.edu/faculty/song/dolap/dolap03/.../p56-Abello.ppt - Similar

LNCS 3255 - Implementing a Query Language for Context-Dependent ... In this paper, we give a short overview of Multidimensional Query Language Those transformations and the MCL query translation are depicted in Figure ... www.springerlink.com/index/21vkrtur/3brcvak0.pdf - Similar by Y Stavraksa - Cited by 9 - Related articles - All 9 versions

1 2 3 4 5 6 7 8 9 10 Next

query translation multidimensional

Sea

Search within results - Language Tools - Search Help - Dissatisfied? Help us im

Google Home - Advertising Programs - Business Solutions - Priva

Web Images Videos Maps News Shopping Gmail more .

Search settings | Sign in



relational query translation multidimensional

Search Advanced Search

Web Show options ...

Results 1 - 10 of about 19,200 for re

Tip: Save time by hitting the return key instead of clicking on "search"

Scholarly articles for relational query translation multidimensional A query language for multidimensional arrays: Design, - Libkin - Cited by 91 nD-SQL: A Multi-dimensional Language for ... - Gingras - Cited by 44 Star/snow-flake schema driven object-relational data ... - Gopalkrishnan - Cited by 25

Relational query processing, query optimization bibliography ... (query processing and optimization); Dynamic multidimensional histograms. (query processing); Rule-Based Translation of Relational Queries into ... scratchpad.wikia.com/.../Relational query processing, query optimization bibliography -Cached - Similar

Modeling multidimensional data sources invention In some embodiments, the facility uses additional relational/multidimensional equivalency logic in its translation of queries and query results. ... www.freshpatents.com/Modeling-multidimensional-data-sourcesdt20070906ptan20070208721.php - Cached - Similar

Object oriented query path expression to relational outer join ... A translation of the object oriented query to a relational query is then density metadata to process multi-dimensional data - A computer readable ... www.freshpatents.com/Object-oriented-query-path-expression-to-relational-outer-jointranslator-method-system-and-article-of-manufact... - Cached - Similar

Show more results from www.freshpatents.com

LNCS 3268 - RAM: A Multidimensional Array DBMS a generic relational kernel. The front-end does not translate the array comprehensions directly into the relational query language of the backend; queries ... www.springerlink.com/index/2tdwv88rr5m8hl11.pdf - Similar by AR van Ballegooij - Cited by 4 - Related articles - All 5 versions

Two-Dimensional Specification of Universal Quantification in a ... Algorithm 2 (Simple-Translation), Input: A two-dimensional query without NE-boxes. Output: A tuple relational calculus query. Constructing a relational ... doi.ieeecomputersociety.org/10.1109/32.126770 - Similar by Q Language - 1992 - All 4 versions

Scientific Commons: Jeffrey F. Naughton

XML Views as Integrity Constraints and their Use in Query Translation (2009) ... Abstract An Array-Based Algorithm for Simultaneous Multidimensional Aggregates Relational query optimizers have traditionally relied upon table ...

en.scientificcommons.org/jeffrey f naughton - Cached - Similar

Multidimensional data clustering scheme for query processing and ... Patent Abstract: Multidimensional clustered tables are provided for efficient ... clustering scheme for query processing and maintenance in relational databases ... or translation707/205, File allocation707/3, Query processing (i.e., ...

www.patentstorm.us/patents/7080081.html - Similar

University of Wisconsin-Madison Database Group Publications Shivnath Babu and Pedro Bizarro: Adaptive Query Processing in the Looking Glass. ... Recursive XML Queries and Relational Storage: XML-to-SQL Query Translation. of Multi-Dimensional Queries in Object-Relational Databases Systems. ... www-db.cs wisc edu/oblications html - Cached - Similar

[PDF] Query Language

File Format: PDF/Adobe Acrobat

Translation of an Existentially Quantified. Two-Dimensional Query to a Relational. Calculus Query. In Section V we discussed how a universally quantified ... ghs1a.home.att.net/publications/TwoDimemsionalUQ.pdf - Similar by KY Whang - 1992 - Cited by 28 - Related articles - All 10 versions

DYNAMIC FILTERS FOR RELATIONAL QUERY PROCESSING _ Patents Patent title: DYNAMIC FILTERS FOR RELATIONAL QUERY PROCESSING ... associated databases can organize and store data in a plurality of multidimensional tables, Patents in class Query formulation, input preparation, or translation ... www.facs.org/aptents/fap/200802/15556 - Cached - Similar

1 2 3 4 5 6 7 8 9 10 Next

relational query translation multidimensional

Search within results - Language Tools - Search Help - Dissatisfied? Help us im

Google Home - Advertising Programs - Business Solutions - Priva

Web Images Videos Maps News Shopping Gmail more .

Search settings | Sign in



relational query translation multidimensional 707/4

Search

Advanced Search

Web Show options...

Results 1 - 10 of about 820 for relation

Relational database management of multi-dimensional data - US ... A relational schema is defined based on the defined multi-dimensional database. ... interface)707/3, Query processing (i.e., searching)707/4, Query formulation, input preparation, or translation707/100DATABASE SCHEMA OR DATA STRUCTURE ... www.patentstorn.us/batents/8205447.hml - Similar

Multidimensional data clustering scheme for query processing and ...
Patent Abstract: Multidimensional clustered tables are provided for efficient ... scheme for query processing and maintenance in relational databases ... 707/100, DATABASE SCHEMA OR DATA STRUCTURE707/4, Query formulation, input preparation, or translation707/20, File allocation707/3, Query processing (i.e., ... www.patentsform.us/patents/7080081.html - Similar

Show more results from www.patentstorm.us

Mapping and query translation between xml, objects, and relations USPTO Application #: 20090150307 - Class: 707 4 (USPTO) ... A single query may access both relational and XML data ... A request is received to perform a query of a multidimensional database, such as a cube, stored at a server computer ... www.freshoatents.com/-d1200906110tan20090150367 onb - Cached - Similar

Extensible command trees for entity data model platform invention USPTO Application #: 20080319957 - Class: 704 (USPTO)... of the query, which further enables translation into Structured Query Language (SQL) and/or ... Subsequently, the nodes can represent various relational and Entity constructs and ... A request is received to perform a query of a multi-dimensional database, ... www.freshpatents.com/-d120081225ptan20080319957.php - Cached - Similar

DYNAMIC FILTERS FOR RELATIONAL QUERY PROCESSING - Patents Origin: CLEVELAND, OH US IPC8 Class: AG06F1730F1 USPC Class: 707 4 ... Relational OLAP (ROLAP), Multidimensional OLAP (MOLAP), and Hybrid OLAP (HOLAP)) are Patents in class Query formulation, input preparation, or translation ... www.fags.org/patents/app/20080215556 - Cached - Similar

Method For Generating A Representation Of A Query - Patent Origin: ANNAPOLIS, MD US IPC8 Class: AG06F1730F1 USPC Class: 707 4 A multidimensional database uses the idea of a data cube to represent For example, the data source may be a relational database or any other type of database. ... Patents in class Query formulation, input preparation, or translation ... www.fags.org/patents/app/20080294612 - Cached - Similar

Show more results from www.fags.org

Mapping of an RDBMS schema onto a multidimensional data model ... Feb 20, 2007 ... 707/4. 5717924, Method and apparatus for modifying existing ... Some relational or multidimensional Databases possess so-called ... dependence suitable for translation into a multidimensional OLAP model The RDBMS Query Technique (QT) is responsible for accessing metadata through a Universe www.freepatentsonline.com/7181440.html - Similar by JCras - 2007 - Cited by 5 - Related articles - All 4 versions

Method of organizing hierarchical data in a relational database ... 5822751, Efficient multidimensional data aggregation operator implementation ... 70714 6085188, Method of hierarchical LDAP searching with relational tables, Bachmann et al. ... Chung et al., A Relational Query Language Interface to a Hierarchical ... Meng et al., A Theory of Translation From Relational Queries to ... www.freepatentsonline.com/6480857.html - Similar by D Chandler - 2002 - Cled by 1 - Related articles - All 2 versions

Show more results from www.freepatentsonline.com

Class Definition for Class 707 - DATA PROCESSING: DATABASE AND ...
Subject matter directed to methods of access, including query path traversal, mapping, and
reuse, joining tables in relational databases, view composition, index choice, E17.056
Multidimensional databases (EPO): ... E17.07 · Query translation (EPO): This subclass is
indented under subclass E17.059 ...
www.uspto.gov/doc/lassification/uspc707/defs707.htm - Cached - Similar

System and method for query planning for execution using ...
Class at Publication: 707/4. Intern'l Class: G06F 17/30 20060101 G06F017/30 In this example, the query framework system 10 has relational planner ...
www.patents.com/System-query-planning...using../en-US/ - Cached - Similar

1 2 3 4 Next

relational guery translation multidimensional 707/4

Search within results - Language Tools - Search Help - Dissatisfied? Help us im

Google Home - Advertising Programs - Business Solutions - Priva